



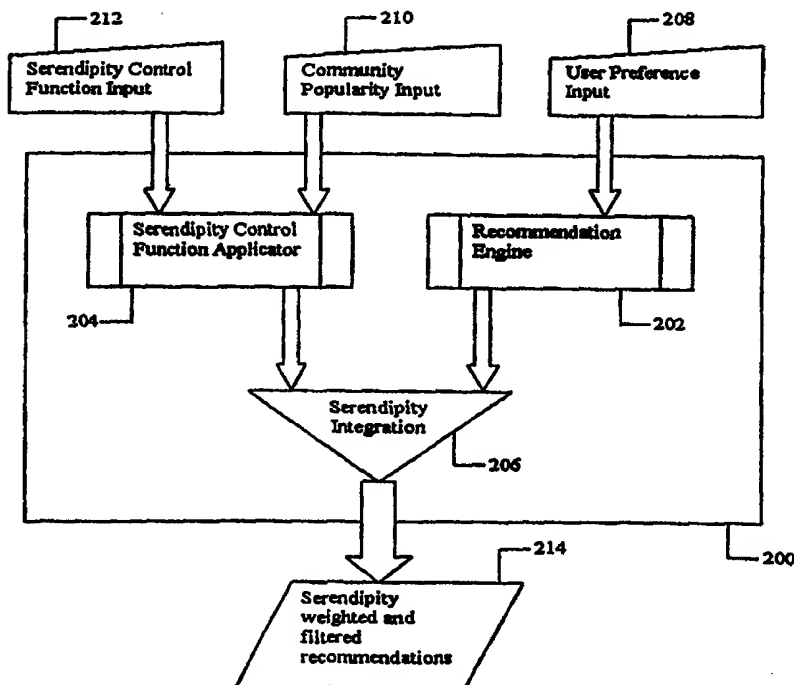
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 17/60, 17/30		A1	(11) International Publication Number: WO 00/04474
			(43) International Publication Date: 27 January 2000 (27.01.00)
(21) International Application Number: PCT/US99/15350 (22) International Filing Date: 7 July 1999 (07.07.99) (30) Priority Data: 09/118,026 17 July 1998 (17.07.98) US (71) Applicant: NET PERCEPTIONS, INC. [US/US]; Suite 300, 11200 West 78th Street, Minneapolis, MN 55344-3814 (US). (72) Inventors: BIEGANSKI, Paul; 6461 Regency Lane, Minneapo- lis, MN 55344 (US). KONSTAN, Joseph, A.; 582 Cretin Avenue South, St. Paul, MN 55116 (US). RIEDL, John, T.; 1340 California Avenue West, Falcon Heights, MN 55108 (US). (74) Agent: DAIGNAULT, Ronald, A.; Merchant & Gould P.C., 3100 Northwest Center, 90 South Seventh Street, Minneapolis, MN 55402-4131 (US).		(81) Designated States: AE, AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the</i> <i>claims and to be republished in the event of the receipt of</i> <i>amendments.</i>	

(54) Title: SYSTEM, METHOD AND ARTICLE OF MANUFACTURE FOR MAKING HIGH USER VALUE RECOMMENDATIONS

(57) Abstract

The invention includes an electronic processing system, a method and a computer readable storage device for generating a serendipity-weighted recommendation output set to a user based, at least in part, on a serendipity function. The system includes a processing system to receive user item preference data and community item popularity data. The processing system is also configured to produce an item recommendation set from the user item preference data, produce a set of item serendipity control values in response to the serendipity function and the community item popularity data, and combine the item recommendation set with the set of item serendipity control values to produce a serendipity-weighted and filtered recommendation output set. The method includes receiving item preference data and community item popularity data. The method further includes producing an item recommendation set from the user item preference data, using the processing system, and generating a set of item serendipity control values in response to the community item popularity data and a serendipity function, also using the processing system. The method also includes combining the item recommendation set and the set of item serendipity control values to produce a serendipity-weighted and filtered item recommendation output set, using the processing system. The computer readable storage device, has a set of program instructions physically embodied thereon, executable by a computer, to perform a method similar to that just described.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

SYSTEM, METHOD AND ARTICLE OF MANUFACTURE FOR MAKING HIGH USER VALUE RECOMMENDATIONS

Background

5 The present invention relates generally to data processing systems, and more particularly, collaborative filtering and recommender systems.

 Recommender systems predict the preferences of users based on attributes known about the user or a past history of preferences or consumption by the user. For example, a recommender system may predict that a user will like the
10 movie “Titanic” because he previously indicated a liking for such other epic movies as “Lawrence of Arabia” or “Ben Hur”.

 It is important for recommender systems to consider quality as well as content of items being recommended. However, there exists no reliable computerized process for estimating the quality of the content of most items that are
15 the subject of a recommender system, especially non-textual multimedia content such as movies and videos. For example, a computer program cannot tell if a remade version of the movie Casablanca is well done or not? Collaborative filtering technology allows recommender systems to provide recommendations based on both
20 content and quality by making recommendations using quality and content judgments made by other users, not by automated computer analysis of content. By pooling together the preferences evinced by a community of users, collaborative filtering allows that user community to share the preferences anonymously in a large scale manner. A prediction of one user’s preference for an item is computed by
25 considering other people’s preferences for that item, where those other people are chosen based on how similar their interests and expectations are to the user’s. For example, a user might be given a high prediction for the movie “Sense and Sensibility,” because other users in the community who shared the user’s taste in Jane Austen movies thought “Sense and Sensibility” was a good movie.

 A recommender system determines its recommendations by
30 examining previous user preference data. The preference data can be unary or numerically valued. Unary preference data is a set of customer-item pairs: a

customer-item pair indicates that an event linking the customer to the item has occurred. No additional preference information is available to the recommender system about the a user-item event except that it happened. The non-existence of a customer-item pair (more generally known as a tuple) for a specific customer-item pair does not indicate a preference: it only indicates a lack of information. An example of unary customer data is purchase record data where a customer-item pair indicates that the customer has purchased the indicated item. Another example of unary data is contained in web page logs, where a customer-item pair indicates that the customer has visited a specific web page.

Binary and numerically valued preference data are generally in the form of a 3-tuples, where the three elements of the tuple are customer identifier, item identifier, and preference value. The preference value indicates, for example, the strength of the user's preference for the item or whether the user's preference is either for or against the item. To illustrate, where the preference is represented in binary form, a "0" may mean a preference against an item while a "1" means a preference for the item. Where the preference is presented as numerically valued data, the data value may represent a one-dimensional axis of preference, with the midpoint indicating an ambivalent preference for the item, a low value indicating a strong dislike for the item, and a high value indicating a strong preference for the item.

Preference data may be presented to the recommender system in explicit or implicit form. Explicit preference data are preference values that a user has supplied directly, for example by filling out a survey. Implicit preference data consist of preference values that have been inferred by observing actions that the user has taken. It can be inferred that the user has some preference for the item that she has just bought, although the act of purchasing the item is not an explicit statement of preference *per se*. A user's preference for a web page may be inferred, for example, by measuring the amount of time that the user spends reading the web page, or the number of times the user returns to that page.

The inputs to a recommender system are typically preference values as described above. The outputs of the recommender system are predictions of preference values for items, particularly those for which the user has not already

indicated a preference. Like the input values, the output preferences may be unary, binary, or numerically valued. A system that outputs unary recommendations predicts items that will be of interest to the user, but does not attempt to predict the strength of a user's preference for each item. Binary predictions indicate items that are likely to be of high preference to the user and items that are likely to be of low preference, but again cannot provide an estimate of preference strength. Numerically valued preferences indicate a preference for or against the item and also indicate the preference strength. Note that the domain of the preference input may be different from the domain of the output preference predictions. For example, the preference input may be unary, while the output preference predictions may be numerically valued.

While unary and binary preference values do not indicate the strength of the preference, some recommender systems may additionally rank the preference predictions being returned such that the highest rank predictions have the largest probability of being correct. Numerically valued items are implicitly ranked.

Existing recommender systems generate recommendations by selecting the highest-ranking positive preference values. However, this technique does not always provide a desirable effect. In many cases, if the recommender system has sufficient data to have high confidence that a recommendation will be good, then the recommendation will be obvious to the user. If the recommendation is obvious, then the recommender system has provided no value. For example, in the context of a system that recommends items to a user for purchase at a grocery store, the recommender system may determine that it is most likely that the user will be interested in purchasing milk, and therefore will recommend that the customer purchase milk. However, a large fraction of grocery store shoppers buy milk, even without being recommended to do so, and so, making a recommendation to purchase milk is obvious. Thus, making such a recommendation to the customer is not very helpful. In spite of being an accurate recommendation, it is not a useful, or valuable, recommendation, since it does not provide the customer with knowledge that he did not already have.

Therefore, there exists a problem with existing recommender systems that, although able to recommend items with high confidence level, often

recommend items that are obvious to the user. Consequently, the value of the recommendation is low. There exists a need to overcome the problem of making low value recommendations.

5

Summary of the Invention

To address the problems listed above, the present invention is directed to an electronic processing system for generating a serendipity-weighted recommendation output set to a user based, at least in part, on a serendipity function. The system includes a processing system of one or more processors configured to
10 receive applicable data that includes user item preference data and community item popularity data. The processing system is also configured to produce an item recommendation set from the user item preference data, produce a set of item serendipity control values in response to the serendipity function and the community item popularity data, and combine the item recommendation set with the set of item
15 serendipity control values to produce a serendipity-weighted and filtered recommendation output set.

In another embodiment, the invention is directed to a method of producing a serendipity-weighted recommendation to a user, the method using a computer having a memory unit, a processing system having one or more processors
20 and an input/output interface. The method includes receiving applicable data by the processing system. The applicable data includes user item preference data and community item popularity data. The method further includes producing an item recommendation set from the user item preference data, using the processing system, and generating a set of item serendipity control values in response to the community
25 item popularity data and a serendipity function, also using the processing system. The method also includes combining the item recommendation set and the set of item serendipity control values to produce a serendipity-weighted and filtered item recommendation output set, using the processing system.

Another embodiment of the invention is directed to a computer-
30 readable program storage device, having a set of program instructions physically embodied thereon, executable by a computer, to perform a method for providing a serendipity-weighted and filtered recommendation. The method includes receiving

applicable data that includes user item preference data and community item popularity data. The method further includes producing an item recommendation set from the user item preference data, generating a set of item serendipity control values in response to the community item popularity data and a serendipity function, and combining the item recommendation set and the set of item serendipity control values to produce a serendipity-weighted and filtered item recommendation output set.

The above summary of the present invention is not intended to describe each illustrated embodiment or every implementation of the present invention. Other features of the invention, together with a fuller understanding of the invention will become apparent and appreciated by referring to the following description and claims taken in conjunction with the accompanying drawings.

Brief Description of the Drawings

The invention may be more completely understood in consideration of the following detailed description of various embodiments of the invention in connection with the accompanying drawings, in which:

FIG. 1 illustrates a computer system for use with an embodiment of the present invention;

FIG. 2 illustrates a system for generating serendipity-weighted and filtered recommendations to a user according to one embodiment of the present invention;

FIG. 3 illustrates a system for generating serendipity-weighted and filtered recommendations to a user according to another embodiment of the present invention;

FIG. 4 illustrates an example of a universe of users, including a customer, and user item preference profiles;

FIG. 5 illustrates a flow chart for one embodiment of a procedure for constructing a neighborhood;

FIG. 6 illustrates a number of different serendipity control functions;

FIG. 7 illustrates a flow chart for one embodiment of a method for calculating item occurrence frequency in a neighborhood;

FIG. 8 illustrates an embodiment of an electronic commerce server for use with the invention;

FIG. 9 illustrates an embodiment of a communications system for a phone call center for use with the present invention;

5 FIG. 10 illustrates one embodiment of a method for calculating community item popularity data;

FIGs. 11 and 12 illustrate embodiments of methods for applying serendipity control functions to community item popularity data to produce serendipity control values;

10 FIGs. 13 and 14 illustrate embodiments of method of integrating serendipity control values with item recommendations; and

FIGs. 15, 16A and 16B illustrate different embodiments of processing systems configured to carry out the invention

While the invention is amenable to various modifications and
15 alternative forms, specifics thereof have been shown by way of example in the drawings and will be described in detail. It should be understood, however, that the intention is not to limit the invention to the particular embodiments described. On the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended
20 claims.

Detailed Description

Present recommenders focus on making accurate recommendations. However, the most accurate recommendations may be those recommendations that are obvious, having low value to the user. For example, a recommender having
25 access to data regarding a user's preference for books may have realized that the user enjoys high-tech military dramas. Therefore, it would be obvious to recommend that the user read a new release by Tom Clancy. However, such a recommendation would most likely be useless to the user, since Clancy is a very popular author and there is a high probability that the user already knew of the recent Clancy release
30 independently of the recommendation. In this case, the recommender adds little value to a user's book selection or to a (virtual) bookstore's ability to serve the user.

The recommender typically has higher value for both the user and the service provider if the item recommended to the user is not known to the user, but is still liked by the user. Thus, a desirable characteristic of recommender systems is known as *serendipity*. Serendipity occurs when a user makes a pleasant discovery that would not normally have been made. A serendipitous recommender system is a recommender system that has the characteristic of providing recommendations that a user would not normally have come across. However, since existing recommender systems focus on providing accurate recommendations, the user is often presented with an obvious recommendation, and not a serendipitous one.

The invention described and claimed below is directed to producing serendipitous recommendations, and may operate under real-time to provide fast interaction with the user or customer. The invention is directed to a recommender system that receives data indicative of a user's preferences, both explicit and implicit preference measures, and generates a serendipity-weighted and filtered recommendation set in response.

FIG. 1 shows a diagram of a typical computer system suitable for practicing the present invention. The computer may include a system of one or more central processing units (CPUs) 102, a memory system 104, an input/output (I/O) adapter 106, a secondary storage 108, a network interface 110, a user interface adapter 114, and a display adapter 112. All of the computer components are connected by a system bus 115. The display adapter 112 may be connected to a display 116 for displaying a recommendation to a user. The user interface adapter 114 may be connected to a user input device 118.

The computer system may include more than one processor, where the processors are in different locations. In such a case, the processors may be linked by input/output interfaces over a network, such as a local area network, wide area network or the Internet.

In one particular embodiment of the invention, recommendations are generated based on user item preference data, population preference data that is a summarization of the preference data of all users, and a serendipity control function, as illustrated in FIG. 2. In this embodiment, the recommender system includes three core components: the recommendation engine 202, the serendipity control function

applicator 204, and the serendipity integration unit 206. The outputs of the control function applicator 204 and the recommendation engine 202 feed into the serendipity integration unit 206. The output from the serendipity integration unit 206 is a serendipity-weighted and filtered recommendation for an item. The
5 serendipity-weighted and filtered recommendation may be generated automatically upon receipt of the user item preference data, or may be generated in response to a specific request received from the user, by way of the user input device 118.

The recommendation engine 202 predicts a user's preference for each item in the dataset, without regard to serendipity. The engine builds and maintains
10 an interest profile for each user that occurs in the user preference data. The recommendation engine can then generate recommendations for each user that occurs in the preference data. There are many different techniques for generating personalized recommendations from user interest profiles that the recommendation engine may use. One example of two techniques that might be used to generate
15 recommendations are automated collaborative filtering as described in Resnick Iacovo, Susha, Bergstrom, and Riedl, GroupLens: An open architecture for collaborative filtering of netnews, Proceedings of the 1994 Computer Supported Collaborative Work Conference (1994).

Other recommendation techniques are described in U.S. application
20 serial no. 08/725,580, filed October 7, 1996, entitled "System, Method And Article Of Manufacture For Utilizing Implicit Ratings In Prediction Information Systems", US application serial no. 08/729,787, filed October 8, 1996, entitled "System, Method, And Article Of Manufacture For Generating Implicit Ratings Based On Receiver Operating Curves" and US application serial no. 08/733,806, filed on
25 October 18, 1996, entitled "System, Method And Article Of Manufacture For Using Receiver Operating Curves To Evaluate Predictive Utility", all of which have been incorporated herein by reference.

In automated collaborative filtering, candidate items for
recommendation are generated by matching users who have shared interests in the
30 past into groups. These groups are called affinity groups or neighborhoods. Members of a user's affinity group are called neighbors. To form a neighborhood for a user, the recommendation engine finds the set of people in the preference data

who have the profiles most similar to the profile of a user. Similarity between two profiles may be measured by counting the items that are shared by the two profiles.

There are many different methods to form affinity groups, and these methods have been described in the literature published on recommendation engines.

5 In the following paragraphs, one particular method is described that may be used to form affinity groups of users based on user preference data. This is provided only as an example and is not intended to limit the invention to the use of the algorithm specifically described. Other examples of neighborhood formation algorithms may be found in Resnick Iacovo, Sussha, Bergstrom, and Riedl, GroupLens: An open
10 architecture for collaborative filtering of netnews, Proceedings of the 1994 Computer Supported Collaborative Work Conference (1994), and Shardanand and Maes Social information filtering: algorithms for automating "word of mouth", CHI '95: Conference proceedings on Human factors in computing systems (1995), both incorporated herein by reference. The recommendation may be made upon receipt
15 of unary, binary, or numerically valued information from the user.

FIG. 4 illustrates the formation of a neighborhood through a simple example of an interaction with a recommendation engine. In this example, the recommendation engine has preference information for ten users – shown in the user universe 400. For generality, the items being recommended are represented by
20 geometric shapes, but this does not limit the applicability of the system to other types of items selected by the user. The list of items that a user has selected or purchased (hereinafter simply selected) is called a profile and the preference contains profiles for each of the ten users of the system. The profiles are represented as small "shopping baskets" beside each user that contain pictorial representations of
25 the items that the user has selected.

Recommendation

One 402 of the ten users, referred to herein as the customer, requests recommendations from a recommender system. The recommendation may be returned to the customer in real time.

30 The customer 402 has selected three items: a circle, a triangle, and a cross. As a result, the customer's profile 404 contains a circle, a triangle and a cross.

To form a neighborhood, the recommendation engine selects the set of users from the user population who have the profiles most similar to the customer's profile. It is important to limit membership in the neighborhood to those neighbors who can contribute positively to a personalized recommendation for the customer. Since user preference data can often contain information for large numbers of users, it is also important to limit the number of users that are selected as neighbors. Therefore, the formation of neighborhoods is controlled primarily by two parameters, namely i) minimum neighbor similarity and ii) maximum neighborhood size. These parameters may be configured by the administrator of the recommendation engine. Minimum neighbor similarity ensures that all members of a neighbor have a minimum number of items in common with the customer before being included in a neighborhood. Configuring the system for a large maximum neighborhood size provides a larger set of items that can be recommended for each user, while a smaller maximum neighborhood size focuses on including only the best neighbors.

In the present example, the minimum neighborhood similarity is set at two shared items, and the maximum neighborhood size is not set (or is set to be equal to the total number of users) because the number of users in the preference data is still small and does not impact on the performance of the recommendation system. As a result, the neighborhood 406 is formed by selecting all the users whose profiles share at least two items with the customer's profile. Therefore the neighborhood is all users who have either: a circle and a triangle; a circle and a cross; or a triangle and a cross. In this case, five users are selected as members of the neighborhood.

A flowchart shown in FIG. 5 illustrates one embodiment of a procedure to be followed to create a neighborhood. At step 500, the neighborhood criteria are set. For example, this may include setting the minimum number similarity, **n_similar**, to be equal to n_0 ($n_0=2$ in the current example). Also, this may include setting the maximum neighborhood size **n_maximum** at n_{max} . The user item preferences of the customer are compared with the preferences of users stored on the database (known as database users), at step 502. The number of similar preferences between the customer's preferences and the database user's preferences are counted, at step 504, and if the number is equal to or higher than **n_similar**, then that particular database user is entered into the neighborhood, step

506. If there are more database users whose preferences have to be examined, as determined at step 508, then the next database user's preferences are compared with the preferences of the customer, at step 510.

Once the preferences of all the database users have been compared
5 against the preferences of the customer, the neighborhood is examined for size, at step 512. If it is determined that the neighborhood is too large, i.e. that the number of members in the neighborhood exceeds **n_maximum**, then the **n_maximum** database users whose preferences are most similar to the preferences of the customer are selected for the neighborhood, in step 514. The neighborhood is now complete,
10 step 516.

Once the neighborhood has been formed, a set of candidate items for recommendation is generated from the neighborhood. The set of candidate items includes all items that occur in profiles of the neighbors but do not occur in the profile of the customer. Eliminating items occurring in the user's profile prevents
15 the recommender system from recommending items that the customer has already selected. The recommendation engine calculates the frequency of occurrence (f_n) of each item within the user's neighborhood. The value of f_n for a particular item is that number of neighbor profiles in which the item occurs relative to the number of neighbors in the neighborhood. The recommendation engine outputs the list of
20 candidate items along with their respective frequency of occurrence in the neighborhood.

A flow chart shown in FIG. 7 illustrates a simple procedure for determining f_n for each item. First, all the items in the neighborhood are selected, and then those items for which the customer has already shown a preference are
25 removed from the list of items, step 702, leaving only those items in the neighborhood that the customer has not selected. For each item, the number of neighbor profiles containing that item is counted, at step 704, to give the item frequency, f_i . It is assumed that there is a total of **total_neigh** neighbors profiled in the neighborhood. The item frequency, f_n , is therefore given by $f_n = f_i / \text{total_neigh}$

30 In the illustrated example, circles, triangles, and crosses are not considered for recommendation because they already occur in the customer's profile. However, hearts, diamonds, and crescents occur in the profiles of neighbors but not

in the customer's profile. The frequency of occurrence for each is computed and output by the recommendation engine. The computed frequencies for the example are shown in Table I. Note that squares are not considered for recommendation because no squares appear in the customer's neighborhood.

5

TABLE I – NEIGHBORHOOD FREQUENCY

Item	Frequency of occurrence (f_n)
Heart	80%
Diamond	60%
Crescent	20%

Community Item Popularity

The community popularity data 210 contains data summarizing the overall popularity of different items selected by members of the population 400. The community popularity data 210 may be stored on a memory system and retrieved when a set of serendipity-weighted and filtered recommendations is being calculated. The community popularity data 210 may also be calculated in real-time upon receipt of a request for the data. The community popularity data 210 may then be calculated from the universe of user preference data, as described in the following paragraphs.

A flow chart, illustrating one embodiment of a procedure for calculating the community popularity data, is illustrated in FIG. 10. Unlike the recommendation techniques described above, the population of interest in calculating the community popularity data is the whole universe of user profiles, and not only those user profiles which demonstrate a similarity to the customer's profile.

The total number of occurrences, f_{univ} , in the universe of user profiles is calculated for each item, in step 1000. The total number of user profiles, T_{univ} , is calculated in step 1010. For each item, the community item popularity, f_p , is given by $f_p = f_{univ}/T_{univ}$.

25

In the illustrated example, 10 users have manifested preferences. A subset of eight of the users have chosen hearts, another subset of six have chosen circles, another subset of five have chosen triangles, another subset of four have chosen crosses, another subset of four have chosen diamonds, another subset of two users have chosen squares, and a last subset of one user has chosen a crescent. In this example, the community popularity of an item is described in terms of the frequency of occurrence of that item in all users' profiles. The community popularity data for the present example may be summarized as shown in Table II.

TABLE II - COMMUNITY POPULARITY INPUT

Item	Popularity (f_p)
Heart	80%
Circle	60%
Triangle	50%
Cross	40%
Diamond	40%
Square	20%
Crescent	10%

Serendipity Control Function

For each item in the community popularity input 210, the serendipity control function applicator 204 computes a serendipity control value by applying a serendipity control function that was input in 212 to each item occurring in community popularity data set 210.

The serendipity control function 212 specifies how community popularity input values 210 are mapped to serendipity control values output from the applicator 204. As just discussed, a community popularity value exists for each item that occurs in the user preference data. In this embodiment, the serendipity control function 212 includes a function that can be specified by the user or by the administrator of the recommender system. The control function may be different for

any (user, item) combination, or may be shared over multiple users and items. In the illustrated example, a heart has a community popularity value of 80%.

The serendipity control function may also be received from an external source, for example the system where the community item popularity database is stored, or may be stored in the memory 104, in the secondary storage 108, or selected by the CPU 102 from a number of functions, based on an input received from the user or a system administrator through the user interface adapter 114. The user may direct a request through the user interface adapter 114

The result of applying the serendipity control function to the community popularity value is the production of a serendipity control value for the item, which is a real-numbered value between 0 and 1, with 1 indicating that the item has met the serendipity requirements of the serendipity control function perfectly, and a 0 indicating that the item does not meet the serendipity requirements of the serendipity control function.

The serendipity control function may be any general mathematical function that relates the community popularity, f_p , of an item to a probability value, w , where $0 \leq w \leq 1$. Therefore, where an item occurs with a specific community popularity value, it is assigned a serendipity control value that is dependent on the community popularity value. Where the serendipity control function is **seren_fun** (**x**), then the serendipity control value, **w**, for a specific item is given by:

$$w = \text{seren_fun}(f_p)$$

Without limiting the scope of the invention, some common strategies for building serendipity control functions are described as follows, with particular reference to FIG. 6.

High frequency cutoff. A high frequency cutoff control assigns a 0 to all items with a community popularity value greater than a specified cutoff value, and a 1 to all other items. A graph of a high frequency cutoff control function is shown in 602. Items that are very popular among the entire user base, and hence have a high community popularity value, are likely to be known by the user.

Recommending an obvious item is not a useful functionality for a recommender system. The high frequency cutoff function provides a method to prevent popular items from being recommended.

In illustration, consider the community popularity input summarized in Table II. A high frequency cutoff function might be configured to remove from consideration any items that have been purchased by more than 75% of the population, since it is believed that the recommendation of items having a community popularity value higher than 75% does not provide any value to the customer. As a result, the high frequency cutoff function assigns a serendipity control value of 0 to any items with f_p greater than 75%, and a 1 to all other items. In this case, the heart is assigned a value of 0, and all other items are assigned a value of 1.

10 This function may be calculated using the following:

if $f_p < 0.75$

then $w = 1$

else $w=0$

15 *Low frequency cutoff.* A low frequency cutoff function is shown in 604. Items with a community popularity value lower than a specified cutoff value are assigned a value of 0, while all remaining items are assigned a value of 1. This function may be useful in cases where it is desirable not to recommend an item until a certain number of users in the population have selected it. An assumption implicit in this function is that the item has some value *per se*, and is therefore worth
20 recommending to a user, only if it has been previously selected at least a given minimum number of times.

Again consider the data illustrated in Table II. In the case where a low frequency cutoff function is chosen so that items with f_p less than 15% are discarded, the crescent is assigned a serendipity control value of 0, since its
25 community popularity value is 10%. All other items are assigned a value of 1.

This function may be calculated using the following:

if $f_p > 0.15$

then $w = 1$

else $w = 0$

30 *High Frequency Dampening.* This function is illustrated as graph 606, and may be described as a broader variation of the high frequency cutoff function. Like the high frequency cutoff function, this control function assigns a

serendipity control value of 1 to items occurring with a community popularity value up to a specified cutoff value. For items occurring with a community popularity value above the specified cutoff value, the control function returns an increasingly smaller value. This function may be used where it is believed that the

5 recommendation of an item is less serendipitous where the item is more popular. However, this function is more forgiving of higher frequency occurrences than the high frequency cutoff function, in that it still allows some items of high community popularity to be assigned a non-zero serendipity control value, albeit at a value less than 1.

10 This function may be calculated using the following:

if $f_p < f_h$

then $w = 1$

else $w = 1 - (f_p - f_h)/m1$

15 where f_h represents the community item popularity above which the cut-off commences, and $m1$ determines the slope of the cut-off. It will be understood that of the cut-offs may be employed other than the linear relationship employed in this example.

Frequency Normalization. In this function, a smaller serendipity control value is associated with an item whose occurrence frequency is greater. The control value decreases linearly as the community popularity of the item increases. An example of a frequency normalization function graph is illustrated as curve 608. The frequency normalization curve need not be linear, and may have a different shape, for example an exponential or polynomial relationship to the community popularity value. A frequency normalization function has the advantage that no matter what the community popularity values of two items are, the less popular item will have a stronger value than the more popular item. Also, the frequency normalization function does not eliminate any item entirely from consideration and does not require the specification of defined cutoff values, which may often be different for different sets of data.

30 A linear frequency normalization function may be given as:

$$w = 1 - f_p/m2$$

where m_2 determines the slope of the normalization.

Dynamically Adapted. A dynamically adapted serendipity control function for a user does not remain constant with time, but changes in response to feedback received from the user. The feedback may be provided as changes
5 explicitly requested by the user, or as information gained by observing the actions of the user. For example, a user may provide an indication of the accuracy of each recommendation given and such indication may be used to identify the level of serendipity that a user appreciates. This in turn may be used to dynamically generate a serendipity control function personalized to that user. For example, where a
10 dynamically adapted function is based on a high frequency cutoff function, the cutoff popularity value may be dynamically controlled. Consequently, the user may control the level of obviousness of the items that are recommended.

The feedback to control the serendipity function may be implicitly received from the customer. For example, the feedback may include the rate at
15 which the user accepts the serendipity-weighted and filtered recommendations presented to him or her. Where the recommendations are accepted at a low rate, the serendipity function may be altered to produce a greater number of serendipity-weighted and filtered recommendations. In addition, the feedback may be based on the rate at which the customer requests additional items or the rate at which the
20 customer accepts non-recommended items. If the rate at which the customer accepts non-recommended items is high, say higher than the rate at which he accepts recommended items, the serendipity function may be altered to increase the range of items recommended to the customer.

Mixed function. The characteristics of more than one of the above
25 functions may be desired in a recommender system. To this extent, a mixed function may be used that combines several of the above mentioned functions into a single function. For example, consider the combined function 610. This function represents a combination of high-frequency cutoff, low-frequency cutoff, and frequency normalization. If the community popularity value is below the low cutoff or above
30 the high cutoff, then the function will assign a serendipity control value of 0 to the item. If the item has a popularity value between the low cutoff and the high cutoff,

then the frequency normalization function applies, with higher popularity values being assigned lower control values.

It will be appreciated that serendipity control functions, other than those described here, may also be used to reduce the obviousness of items recommended by a recommender system.

A flow chart showing one embodiment of a method of applying the serendipity control function is illustrated in FIG. 11. In this real-time embodiment of a serendipity control function applicator, the serendipity control value is calculated in response to a request received for a serendipity control value of an item, at step 1100. The frequency of occurrence of the item in the universe of user profiles, f_{univ} , and the total number of user profiles T_{univ} are retrieved from a database of user profiles, in step 1102. The community item popularity is determined from f_{univ} and T_{univ} , and the serendipity function applied to the community item popularity value to produce the serendipity control value, at step 1104.

The method illustrated in FIG. 11 to generate the serendipity control value may be carried out in a processor that is remote from the processor that calculates the serendipity-weighted and filtered recommendation. Therefore, the serendipity control value may be output to that processor, at step 1106. However, the same processor may also be used both to calculate the serendipity control value and the serendipity-weighted and filtered recommendation. This serendipity control function applicator operates in real-time, calculating the serendipity control value afresh each time there is a request.

Another method for calculating the serendipity control value is illustrated in flow-chart form in FIG. 12. In this embodiment of a serendipity control function applicator, the serendipity control value is not calculated on demand, as is the case in the embodiment shown in Fig. 11. Instead, the serendipity control value is calculated in advance of such time as it is needed, and is stored on a memory, readily accessible for retrieval. This approach may save processing time compared with the method illustrated in FIG. 11. However, there is a possibility that the community item popularity data used in this method may not be as up-to-date as the community item popularity data calculated in the method of FIG. 11.

In this method, a first item in the user preference database is selected, at step 1200. The frequency of that item's occurrence in the universe of user profiles, **freq-occurrence**, and the total number of user profiles, **num_users**, are retrieved, and the community item popularity for that item is calculated, at step 5 1202. The serendipity control value is then calculated from the community item popularity and the serendipity function, at step 1204, and the serendipity control value is stored in a database, available for lookup at a later time, at step 1206. If it is determined, at step 1208, that there are more items to be considered, then the processor selects the next item, at step 1210, and proceeds to calculate and store a 10 serendipity control value for the next item in a similar manner. Once all items have been considered, the database of serendipity control values is complete, step 1212.

This method is suitable for operating in the background, for example when the processor is not being heavily used. In addition, once the serendipity control value database is complete, the serendipity control function applicator may 15 restart, and calculate a new set of control values. However, for applications where the volume is large, and the frequency of occurrence is not expected to vary significantly over time, the serendipity control value database may be updated at less frequent intervals, perhaps daily or weekly.

Considering once more FIG. 2, the serendipity integration unit 206 20 takes the serendipity control value of an item produced by the control function applicator 204, and the recommendation value of that item produced by recommendation engine 202, and performs a computation to produce a final serendipity weighted recommendation.

Serendipity Integration

25 The purpose of the serendipity integration unit is to reconcile the prediction accuracy of the recommendation engine and the serendipitous value of the recommendation into a final list of weighted and filtered recommendations. The serendipity integration may be any computation that relates a serendipity control value w and a recommendation value r of an item to a serendipity-weighted and 30 filtered recommendation r_w . The recommendations presented to the serendipity integrator may be unary, binary or numerically valued in form. Likewise, the

serendipity-weighted and filtered recommendations may also be unary, binary or numerically valued in form.

Some possible serendipity integration strategies include:

- Product. $r_w = w * r$. The recommendation produced by the recommendation engine, r , is devalued by the serendipity control value. Items that meet the serendipity requirements perfectly (thus having a value of 1) do not devalue the recommendation at all, and $r_w = r$. Items that do not meet the serendipity requirements at all (having a value of 0) result in a weighted recommendation of 0. Recommendations for items with a control value between 0 and 1 are devalued by an amount inversely proportional to the value of the serendipity control value w .

Consider again the example described in FIG. 4, where the further assumption is made that a low frequency cutoff serendipity control function is being used, and that the low cutoff is set to 15%. From Table I, we see that the recommendation engine produces recommendations for the heart, diamond, and crescent, with recommendation values of 80%, 60%, and 20% respectively. From Table II, we see that the heart and the diamond have community popularity values of 80% and 40% respectively, and therefore the function applicator 204 assigns a serendipity control value of 1 to both the heart and the diamond.

The serendipity integration unit computes serendipity weighted recommendations for the heart and the diamond by applying the product, resulting in weighted and filtered recommendations 80% ($1 * 80\%$) and 60% ($1 * 60\%$). The weighted and filtered recommendations in this case are exactly the same as the original recommendations.

However, the crescent has a community popularity value of 10%, which causes the function applicator 204 to assign it a serendipity control value of 0. This causes the integration unit 206 to assign to the crescent a control value of 0, after applying the product. This effectively filters out the crescent from the set of serendipity-weighted and filtered recommendations.

This may further be described mathematically as:

$$r_w = r * w, \text{ where}$$

$$w = 1 \text{ for } f_p \geq 0.15$$

$$w = 0 \text{ for } f_p < 0.15$$

Product with cutoff. The serendipity integration unit applies the product as described above, but if the serendipity-weighted and filtered recommendation computed is below a specified cutoff, then the integration unit does not output a recommendation for that item.

This may further be described mathematically as:

$$r_w = r * w, \text{ where}$$

$$w = 1 \text{ for } f_p \geq 0.15$$

$$w = 0 \text{ for } f_p < 0.15$$

$$r_w \geq r_i \text{ and}$$

$$r_w = 0, \text{ where}$$

$$r_w < r_i$$

with r_i being the low cut-off value for the recommendation.

Product with top n. The serendipity integration unit computes serendipity weighted recommendations for all items, and only outputs recommendations for the items with the n largest weighted recommendations.

Product with cutoff and top n. A combination of the two previously described computations, where first items with weighted recommendations below the cutoff are discarded, and then from the remaining items, the ones with the top n highest weighted recommendations are output. If less than n items remain after discarding those below the cutoff, then weighted recommendations for all the remaining items are output.

Regardless of the computation, the serendipity integration unit outputs a set of serendipity-weighted and filtered recommendations. The set of recommendations may be empty. The serendipity integration unit may sort the weighted and filtered recommendations before returning them to the user, or it may return the weighted recommendations in the order that the recommendation request list was received. The weighted and filtered recommendations are output into a list 214. The list of serendipity-weighted and filtered recommendations may be output through the display adapter 112 for direct display on a display device 116, for example a user's computer screen. The serendipity integrator 206 may also store

the set of serendipity-weighted and filtered recommendations to a memory system, for example a secondary storage system 108 so that the user, or system administrator can have access to the set of serendipity-weighted and filtered recommendations at a later time.

5 FIG. 13 illustrates, in flow chart form, one embodiment of a method of operating a processing system, of one or more processors, to produce a set of serendipity-weighted and filtered recommendations. First, a set of item recommendations are received from a recommendation engine 202 at step 1300. The first item in the set of item recommendations is selected, at step 1302, and a
10 request is made for a serendipity control value for the first item, at step 1304. The serendipity control value may be stored in a serendipity control value database, as described earlier with regard to FIG. 12. Alternatively, the serendipity control value may be calculated in real-time, as described earlier with regard to FIG. 11. Once the serendipity control value has been received or calculated, the serendipity-weighted
15 and filtered recommendation is produced by integrating the item recommendation with the serendipity control value, at step 1306, in this case by multiplying the recommendation value by the serendipity control value, and then it is added to the set of serendipity-weighted and filtered recommendations, at step 1308.

 If it is determined, at step 1310, that there are more item
20 recommendations to be considered, the processing system selects the next candidate item, at step 1312, and proceeds with the integration process as just described. Once all item recommendations have been considered, the set of serendipity-weighted and filtered recommendations is output to the customer.

 Another embodiment of a method for integrating the serendipity
25 control values with the item recommendation set is illustrated in flow-chart form in FIG. 14. In this embodiment, a set of item recommendations is received from a recommendation engine 202 at step 1400. The first item in the set of item recommendations is selected, at step 1402, and a request is made for a serendipity control value for the first item, at step 1404. As described above with regard to step
30 1304 in FIG. 13, the serendipity control value may be calculated and delivered in real time, or it may be retrieved from a serendipity control value database. Next, the weighted recommendation is tested against the cut-off value r_i . If it is determined

that the weighted recommendation is less than the cut-off value, at step 1410, then that item is removed from further consideration. On the other hand, if the weighted recommendation is higher than the cut-off value, then the item is added to the weighted and filtered recommendation list, at step 1412.

5 If it is determined that there are more items to be considered, at step 1414, the next item candidate is considered in turn, at step 1416. If there are no more item recommendations to consider, the weighted and filtered recommendation list may be sorted, at step 1418, into order of descending weighted and filtered recommendation value, and output for presentation to the user, at step 1420.

10 FIG. 3 depicts a second embodiment of the present invention. This embodiment of a serendipity applicator and integrator 300 is useful for application to a set of recommendation input data 306 produced by a separate recommender system, which may be located on a different processing system than the serendipity control function applicator and/or the serendipity integration unit. The output from
15 the applicator/integrator 300 is a set of serendipity-weighted and filtered recommendations. The recommendation input data 306 may be a set of tuples of the form (user, item, preference value), indicating users' expressed or observed preferences for items. The preference value may be unary, binary, or numerically valued. The recommendation input data may come from any source, including
20 another recommender system that does not support serendipitous recommendations. The community popularity input 304 contains the community popularity values for each item in the recommendation input 306. The serendipity control function applicator 204 computes serendipity control values by applying the serendipity control function input in 302 to the community popularity values 304. The
25 serendipity integration unit 206 combines the recommendation input and the serendipity control values and outputs serendipity weighted and/or filtered recommendations 312.

 It should be understood that the present invention is not restricted to operating on a single processor system. For example, in the embodiment illustrated
30 in FIG. 2, the serendipity control function applicator 204 may be located on a different processing system from the serendipity integrator 206, while the recommendation engine 202 may be located on yet another processor. However,

any two of these three system elements, the recommendation engine 202, the function applicator, 204 and the integrator 206 may be on one processor with the remaining one being located on a separate processor, or all three 202, 204 and 206 may be on a single processor. In the embodiment illustrated in FIG. 3, the serendipity control applicator 204 and the serendipity integrator 206 may be on different processors that together form a processor system, or may be on a single processor. The recommendation input 306 may be received from a recommendation engine collocated on a processor with the function applicator 204 and/or the integrator 206.

The serendipitous recommender system and the serendipity applicator/integrator have use in many different applications, some of which are described hereinbelow. In order to implement a serendipitous recommender system or a serendipity applicator/integrator, the system administrator devises an appropriate strategy for the serendipity control function, and has access to user preference data. In the following, we list two specific instances of applications of the invention, and then describe different content areas where the invention is useful.

The present invention can be used in different consumer environments, including those where the number of items available for selection by a user is very large. In such an environment, the consumer typically cannot, or will not, spend sufficient time to browse through all of the items in order to determine which item to select.

As an example of an application, the invention may be used to augment an electronic commerce server on the Internet. FIG. 8 shows an illustration of a recommender system integrated into a web-based commerce site. The customer 802 connects using a network 803 to a web server 804. A commerce server 806, connected to the web server 804, processes all financial transactions for the customer and contains a database of products for sale. The web server 804 presents this set of products for sale to the customer 802. A purchase database 808, coupled to the commerce server 806, contains records of previous purchases by the customer 802 and other customers.

A serendipity-weighted recommender system 810, coupled to the web server 804, the commerce server 806 and the purchase database 808, receives

purchase information from the purchase database 808. The recommender system 810 may also receive explicit ratings from the customer 802 via the web server 804 to predict the customer's preference on recommended items that the customer 802 has not yet purchased. These recommendations may then be used by the web server 804 and the commerce server 806 to provide the customer 802 with specifically targeted content, such as recommendations to purchase specific items, recommendations to view specific items, or targeted advertisements.

As a specific example of a recommender system implemented as described above, consider an web-based book store. The commerce server 806 contains a database of all books available for sale, indexed by ISBN, title, author, and subject. The web server 804 provides a web-based interface that allows the user to locate books of interest in two ways. If the customer is not looking for a specific book, then the web server can present the user with a list of personalized book recommendations, provided by the serendipity-weighted recommender system 810, with the list of recommendations possibly limited to be within a specific subject, or by a specific author. If the customer is looking for a specific book, then the customer may search by ISBN, title, author, or subject, and when the customer locates the book of interest, the customer may be presented with the serendipity-weighted and filtered recommendation value for that book as a prediction for how much she will like the book. The serendipity-weighted recommender system 808 provides serendipity-weighted and filtered recommendations using the past purchase history of the customer and of other customers, found in the purchase database 808. These recommendations may be provided in real-time. However, to get a more detailed information on the likes and dislikes of the customers, the customer may be asked to provide explicit preference information, such as rating books that she has read. An example of this would be asking customers to rate books on a scale of 1 to 5 with 1 being "awful" and 5 being "excellent." This explicit preference information would be collected by the web server 804 and transmitted to the recommender system 810 to be used as preference information for recommendations.

Another application in the generic domain is support for phone call centers. Phone call centers include the computer systems and phone switches used to support operators for organizations that take product orders by telephone. The call

center computer systems record and process all phone orders, and may keep customer records of previous orders. FIG. 9 shows a call center implementation that uses the invention. The customer 902 phones into the call center 904, through a telephone network 903, seeking to purchase an item. The customer 902 is connected to an available call center operator, who has access to a call center console 906. A commerce server 908 processes all purchasing and financial transactions, recording purchase information in a purchase database 910. A serendipity-weighted recommender system 912 uses the purchase information stored in the purchase database 910 to generate serendipity-weighted and filtered recommendations for the customer 902. These recommendations are communicated to the call center operator through the call center console. The call center operator may then ask if the customer 902 would be interested in the recommended items. For example, with a phone-in bookstore, the serendipity-weighted recommender system 912 might indicate that the customer is likely to enjoy a specific new release. In offering the recommended item to the customer 902, the call center operator may also offer a discounted price or some other incentive to the user to make a purchase.

Another application of the invention is to predict serendipitous finds in items having audio/visual content. Such items include television programs, movies, radio programs, or a digital media programs. Preference data may be collected by surveying people to learn what programs they currently enjoy, recording which videos they rent from a video store or, in the case of video-on-demand or pay-per-view, recording which programs the user requests. The recommender system recommends new programs that the user may enjoy. Another multimedia content area is music (CDs, tapes, DVDs, etc.). The users may discover musical artists that they would have never originally learned about. A music store may record all music purchases by customers, and use that information as user item preference information. The music store may then provide serendipitous recommendations for new music albums, perhaps in a newsletter or through an information kiosk located in the store or through a web based interface. Still media such as images, pictures, and works of art could also be recommended. The quality and value to a user of such still media varies greatly from person to person. A personalized recommendation from the preferred recommendation system helps a

user to locate new images, pictures, and works of art in a way that ensures quality, interest, and serendipity.

Another application area is textual information sources, which includes content collections, such as books or magazines, and other publications.

5 Textual information sources may also include individual articles from magazines, newspapers, or other article publishers, and include Internet documents which are large in number, and variable in quality. With the explosion of electronic textual information databases, there are more such information sources than any one person has the time to consider. A serendipity-weighted recommender system according to
10 the present invention may help a user to expand his or her interests by finding new, high quality publications that they find interesting. The serendipitous quality of the recommender system provides a value to the user that a standard content-based recommender system cannot offer.

In comparison, text and bibliography search engines locate
15 documents that contain certain specified search terms or words. Search engines typically cannot differentiate the same words used in different contexts, and have no concept of the "quality" of a document. Likewise, traditional text search engines have no way of determining the serendipitous value of a textual item. However, a standard search engine may be coupled with the invention to provide serendipitous
20 finds. One possible technique is to have the invention order the list of items returned by a search engine by the serendipitous recommendation value. Therefore, the term "item recommendation data" should be understood to include results, not only from a recommender system, but also from a search engine. Also, the term "recommendation engine" should be understood to cover search engines.

25 In the marketplace of consumable goods, which includes grocery items, clothing, stationary supplies, and household supplies, there is a large variety of goods and brands available for most items. There is also a large variety of quality, style, and taste among these goods and brands. A supermarket (virtual or physical) may provide the preferred recommender system as a way to help
30 customers locate goods and brands of the appropriate quality, as well as introduce users to new consumable items that they will enjoy that they might not have known about before.

Many supermarkets now provide customers with “shopping club” cards to track individual customer purchases, and to provide customer rewards for frequent shopping. Each customer presents her card at the checkout counter to receive discounts, and the checkout register records all of the items that the customer has purchased. This purchase record may be used as a personal profile for the generation of serendipitous recommendations. The supermarket may then provide the customer with incentives to buy the recommended items by mailing coupons to the customer for those items identified in the set of serendipity-weighted and filtered recommendations.

10 The invention may also be used to help a user locate services of interest to the user. Many people enjoy variety in dining and entertainment services and are interested in trying new and different restaurants, different meal options at restaurants, new theater shows, and different bars and nightclubs. The invention may be used to enhance the variety of entertainment services using serendipity, 15 while ensuring that only services of interest to the user and of the appropriate quality are recommended. Other services that may be recommended include financial services, real estate, architectural goods or services, travel related goods or services, and automobiles, particularly purchasing, renting and repairing.

Another content domain which may benefit from the invention is 20 computer hardware, software, and hardware/software service products. The hardware/software market is an extremely competitive one. Consequently, there is a large selection of brand alternatives, and a large variance in quality among the brands. Also, many new hardware and software products are introduced each year, often too many for a user to keep track of. The invention has the ability to help a 25 user locate hardware and software brands of interest, as well as locate new hardware and software that the user may not have considered buying before.

Another embodiment of the invention, illustrated in FIG. 15, includes a processing system 1500 of one or more processors configured to receive applicable data including i) item recommendation data 306, and ii) community item popularity 30 data 304. The processing system is configured to produce a set of item serendipity control values 1502 in response to a serendipity function 302 and the community item popularity data 204. The processing system 1500 is also configured to combine

the item recommendation data 306 with the set of item serendipity control values 1502 to produce a serendipity-weighted and filtered recommendation output set 312.

The processing system may include only a single processor, or may include a number of different processors. One embodiment of the processing system is illustrated in FIG. 15, where the processing system includes a first processor 1504 (shown in dashed lines) configured to receive the community popularity data 304 and produce the set of item serendipity control values 1502 in response to the serendipity control function 302 and the community popularity data 304. A second processor 1506 (also in dashed lines) is configured to receive the item recommendation data 306 and to combine the item recommendation data 306 with the set of item serendipity control values 1502 to produce the serendipity-weighted and filtered recommendation output set 312. The second processor may be remote from the first processor and coupled to the first processor through a network, the Internet, or some other communications channel.

Another embodiment of the invention, illustrated in FIG. 16A, includes a processing system 1600 of one or more processors configured to receive applicable data including i) user item preference data 208, and ii) community item popularity data 210. The processing system is configured to produce 202 an item recommendation set 1603 from the user item preference data 208, to produce a set of item serendipity control values 1602 in response to the serendipity function and the community item popularity data, and to combine the item recommendation 1603 set with the set of item serendipity control values 1602 to produce a serendipity-weighted and filtered recommendation output set 214.

The processing system may be formed from a single processor adapted to perform recommendation 202, serendipity function application 204 and serendipity integration 206. The processing system may also be formed from a number of processors, for example a first processor 1604 (dashed lines) may be configured to receive the user preference data 208 and generate the item recommendation set 1603, while a second processor 1606 (dashed lines) is configured to receive the community item popularity data 210 and the serendipity control function 212, and to generate a serendipity control value 1602. A third processor 1608 (dashed lines) is configured to receive the item recommendation set

1603 from the first processor 1604 and the serendipity control value 1602 from the second processor 1606, and to produce the serendipity-weighted and filtered recommendation output set 214. The processors may all be remote from each other and intercoupled through a network, the Internet, or some other communications channel.

The processing system may also be formed from two processors, for example as illustrated in FIG. 16B. Here, the first processor 1608 is configured to receive the user item preference data 208, the community item popularity data 210 and the serendipity control function 212. The first processor is further configured to produce the item recommendation set 1603 from the user item preference data 208, and to produce a set of serendipity control values from the serendipity function 212 and the community item popularity data 210. The second processor 1612 is configured to combine the item recommendation set 1603 with the set of serendipity control values 1602 to produce the set of serendipity-weighted and filtered recommendations 214.

It will be appreciated that the processing system 1600 may also be formed from other combinations of two processors. For example, a first processor may be configured to perform serendipity function application 204 while a second processor is configured to perform recommendation 202 and serendipity integration 206. In addition, a first processor may be configured to perform recommendation 202 while a second processor is configured to perform serendipity function application and serendipity integration 206. Like the embodiments described above, the processors may be remote from each other and interconnected through a network, the Internet or some other communications channel.

As noted above, the present invention is applicable to recommender systems. It is believed to be particularly useful in permitting recommender systems to produce recommendations that are not only accurate but also of high value to the user, i.e. recommendations that may not be obvious. Accordingly, the present invention should not be considered limited to the particular examples described above, but rather should be understood to cover all aspects of the invention as fairly set out in the attached claims. Various modifications, equivalent processes, as well as numerous structures to which the present invention may be applicable will be

readily apparent to those of skill in the art to which the present invention is directed upon review of the present specification. The claims are intended to cover such modifications and devices.

WE CLAIM:

1. An electronic processing system for generating a serendipity-weighted recommendation output set to a user based, at least in part, on a serendipity function, the system comprising a processing system of one or more processors
5 configured to:
 - a. receive applicable data including
 - i. user item preference data, and
 - ii. community item popularity data,
 - b. produce an item recommendation set from the user item
10 preference data;
 - c. produce a set of item serendipity control values in response to the serendipity function and the community item popularity data, and
 - d. combine the item recommendation set with the set of item serendipity control values to produce a serendipity-weighted and filtered
15 recommendation output set.
2. A system as recited in claim 1, wherein the processing system is further configured to
 - a. receive the serendipity control function and
 - b. produce the set of item serendipity control values in response
20 to the received serendipity function and the community item popularity data.
3. A system as recited in claim 1, wherein the processing system is further configured to
 - a. receive an input from one of a serendipity filter system administrator and the user relating to the serendipity control function,
 - b. select the serendipity function in response to the input
25 received from the one of the serendipity filter system administrator and the user, and
 - c. produce the set of item serendipity control values in response to the selected serendipity function and the community item popularity data.

4. A system as recited in claim 1, further comprising an input interface operatively coupled to the processing system in order to receive the applicable data and transmit the applicable data to the processing system.

5. A system as recited in claim 4, wherein

5 a. the applicable data further includes the serendipity function,
b. the input interface is configured to be further coupled to a memory system , and

c. the processing system is configured to
i) receive the serendipity function from the memory system ,
10 and
ii) produce the set of serendipity control values in response to the community item popularity data and the serendipity function received from the memory system.

6. A system as recited in claim 4, wherein

15 a. the input interface is configured to be further coupled to receive a request from the user for serendipity-weighted and filtered recommendations, and

b. the processing system combines the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set in response to the request
20 received from the user.

7. A system as recited in claim 4, wherein

the input interface is further configured to receive user item preference data that includes at least one of unary values, binary values, and numerical values, and
25

the processing system is further configured to produce the item recommendation set from the at least one of the unary values, binary values and numerical values.

8. A system as recited in claim 1, further comprising an output interface operatively coupled to the processing system in order to receive the serendipity-weighted and filtered recommendation output set.

9. A system as recited in claim 8, further comprising a display device
5 operatively coupled to the output interface in order to display the serendipity-weighted recommendation output set.

10. A system as recited in claim 8, further comprising a memory system operatively coupled to the output interface in order to receive and store the serendipity-weighted recommendation output set.

10 11. A system as recited in claim 1, further comprising
a memory system configured to be operatively coupled to
the processing system,
wherein the processing system is configured to receive
community item popularity data from the memory system.

15 12. A system as recited in claim 1, wherein the processing system
includes a single processor configured to
a. receive the applicable data including
i. user item preference data, and
ii. community item popularity data,
20 b. produce the item recommendation set from the user item
preference data;
c. produce the set of item serendipity control values in
response to the serendipity function and the community item popularity data,
and
25 d. combine the item recommendation set with the set of item
serendipity control values to produce the serendipity-weighted and filtered
recommendation output set.

13. A system as recited in claim 1, wherein the processing system includes
- a) a first processor configured to
 - i) receive the community item popularity data, and
 - ii) produce the set of item serendipity control values in response to the serendipity control function and the community popularity data, and
 - b) a second processor, configured to be operatively coupled to the first processor and configured to
 - i) receive the user item preference data,
 - ii) produce the item recommendation set from the user item preference data,
 - iii) receive the set of item serendipity control values from the first processor, and
 - iii) combine the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.
14. A system as recited in claim 1, wherein the processing system includes
- a) a first processor configured to
 - i) receive the community item popularity data, and
 - ii) produce the set of item serendipity control values in response to the serendipity control function and the community popularity data,
 - b) a second processor configured to
 - i) receive the user item preference data, and
 - ii) produce the item recommendation set from the user item preference data, and
 - c) a third processor configured to be operatively coupled to the first and second processors, and configured to

i) receive the set of item serendipity control values from the first processor and the item recommendation set from the second processor, and

5 ii) combine the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.

15. A system as recited in claim 1, wherein the processing system includes

a) a first processor configured to

10 i) receive the user item preference data, and

ii) produce the item recommendation set from the user item preference data, and

b) a second processor configured to

i) receive the community item popularity data, and

15 ii) produce the set of item serendipity control values in response to the serendipity control function and the community popularity data,

iii) receive the item recommendation set from the first processor, and

20 iv) combine the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.

16. A system as recited in claim 1, wherein the processing system is configured to produce the serendipity-weighted recommendation output under real-time, interactive time constraints.

25

17. A system as recited in claim 1, wherein the processing system is further configured to produce the serendipity-weighted and filtered recommendation output set as at least one of unary, unordered recommendations and priority-ordered recommendations.

18. A system as recited in claim 1, wherein the serendipity function is at least one of:

- a) a fixed function having controllable parameters;
- b) a bi-level function that excludes items occurring in the community item popularity data with a frequency greater than a pre-selected upper frequency value;
- c) a bi-level function that excludes items occurring in the community item popularity data with a frequency less than a pre-selected lower frequency value;
- d) a continuous function having a value that reduces with increasing frequency of occurrence in the community item popularity data; and
- e) a function assigning a constant value to items having an occurrence frequency in the community item popularity data less than a selected frequency and a value that reduces with an occurrence frequency higher than the selected frequency.

19. A system as recited in claim 1, wherein

- a. the processing system is further configured to select the serendipity function using feedback data, the feedback data including one of a rate at which the user accepts items from the serendipity-weighted and filtered item recommendation set, a rate at which the user requests additional predictions, and a rate at which the user accepts non-recommended items, and
- b. the processing system produces the set of item serendipity control values in response to the community item popularity data and the serendipity function selected using the feedback data.

20. A system as recited in claim 1, wherein

- a. the item recommendation data relate to one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services,

financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-related goods, travel-related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products, and

b. the serendipity-weighted and filtered recommendation output set relates respectively to the one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services, financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-related goods, travel-related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products.

21. A method of producing a serendipity-weighted recommendation to a user, the method using a computer having a memory unit, a processing system having one or more processors and an input/output interface, the method comprising:

- a. receiving, by the processing system, applicable data including,
 - i. user item preference data and
 - ii community item popularity data;
- b. producing, using the processing system, an item recommendation set from the user item preference data;
- c. generating, using the processing system, a set of item serendipity control values in response to the community item popularity data and a serendipity function; and
- c. combining, using the processing system, the item recommendation set and the set of item serendipity control values to produce a serendipity-weighted and filtered item recommendation output set.

22. A method as recited in claim 21, further comprising

- a. receiving, by the processing system, the serendipity control function, and

b. generating the set of item serendipity control values in response to the serendipity control function received by the processing system and the community item popularity data.

23. A method as recited in claim 21, further comprising

5 a. receiving, by the processing system, an input from one of a serendipity filter system administrator and the user relating to the serendipity control function,

b. selecting, by the processing system, the serendipity function in response to the input received from the one of the serendipity filter system administrator and the user, and

10 c. producing, by the processing system, the set of item serendipity control values in response to the selected serendipity function and the community item popularity data.

24. A method as recited in claim 21, further comprising

15 a. receiving, through the input/output interface, a serendipity function selection control input from one of a recommender system administrator and the user, and

b. selecting, using the processing system, the selectable serendipity function in response to the received serendipity function selection control input.

20

25. A method as recited in claim 21, further comprising receiving, by the input/output interface the applicable data and transmitting the applicable data, by the input/output interface, to the processing system.

26. A method as recited in claim 21, further comprising controlling the serendipity function, using the processing system, with feedback data received through the input/output interface, the feedback data including one of a rate at which the user accepts serendipity-weighted and filtered recommended items, a rate at which the user requests additional serendipity-weighted and filtered recommendations, and a rate at which the user accepts non-recommended items.

25

27. A method as recited in claim 21, further comprising
- a. receiving, by the processing system, the serendipity function from the memory system, and
 - b. generating, by the processing system, the set of serendipity control values in response to the community item popularity data and the serendipity function received from the memory system.
28. A method as recited in claim 27, further comprising
- a. receiving a request, with the processing system, from the user for serendipity-weighted and filtered recommendations, and
 - b. combining, with the processing system, the item recommendation data with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set in response to the request received from the user.
29. A system as recited in claim 21, further comprising
- receiving, with the processing system, user item preference data that includes at least one of unary values, binary values, and numerical values, and
 - producing, with the processing system, the item recommendation set from the at least one of the unary values, binary values and numerical values.
30. A method as recited in claim 21, further comprising outputting, using the processing system, the serendipity-weighted and filtered item recommendation output set to the input/output interface.
31. A method as recited in claim 30, further comprising transmitting, using the input/output interface, the serendipity-weighted and filtered recommendation output set to a display device and displaying the serendipity-weighted and filtered recommendation output set on the display device.

32. A method as recited in claim 30, further comprising transmitting, using processing system, the serendipity-weighted and filtered recommendation output set to the memory system for storage therein.

5 33. A method as recited in claim 21, further comprising
receiving, using the processing system, the community
item popularity data from the memory system.

34. A method as recited in claim 21, further comprising
a) receiving the applicable data, with a single processor of the
one or more processors,
10 b) producing the item recommendation set, with the single
processor, from the user item preference data,
c) generating the set of item serendipity control values, with
the single processor, in response to the serendipity control function
and the community item popularity data, and
15 d) combining, with the single processor, the item
recommendation set with the set of item serendipity control values to
produce the serendipity-weighted and filtered recommendation output
set.

35. A method as recited in claim 21, further comprising
20 a) receiving the community item popularity data with a first
processor of the one or more processors
b) producing, with the first processor, the set of item
serendipity control values in response to the serendipity control
function and the community popularity data,
25 c) receiving the user item preference data with a second
processor of the one or more processors,
d) producing the item recommendation set, with the second
processor, from the user item preference data;

- 5 e) receiving, with the second processor, the set of item serendipity control values from the first processor, and
- f) combining, with the second processor, the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.
36. A method as recited in claim 21, further comprising
- 10 a) receiving the community item popularity data with a first processor of the one or more processors
- b) producing, with the first processor, the set of item serendipity control values in response to the serendipity control function and the community popularity data,
- c) receiving the user item preference data with a second processor of the one or more processors,
- 15 d) producing the item recommendation set, with the second processor, from the user item preference data;
- e) receiving, with a third processor of the one or more processors, the set of item serendipity control values from the first processor and the item recommendation set from the second processor, and
- 20 f) combining, with the third processor, the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.
- 25 37. A method as recited in claim 21, further comprising producing, with the processing system, the serendipity-weighted and filtered item recommendation output set under real-time, interactive constraints.
38. A method as recited in claim 21, further comprising producing, with the processing system, the serendipity-weighted and filtered recommendation output

set as at least one of unary, unordered recommendations and priority-ordered recommendations.

39. A method as recited in claim 21, further comprising receiving, with the processing system, the item recommendation data as at least one of unary values,
5 binary values, and numerical values.

40. A method as recited in claim 21, further comprising generating, using the processing system, the set of item serendipity control values in response to the community item popularity data and the serendipity function, where the serendipity function is at least one of:

- 10 a fixed function having controllable parameters;
- a bi-level function that excludes items occurring in the community item popularity data with a frequency greater than a pre-selected upper frequency value;
- a bi-level function that excludes items occurring in the
15 community item popularity data with a frequency less than a pre-selected lower frequency value;
- a continuous function having a value that reduces with increasing frequency of occurrence in the community item popularity data; and
- 20 a function assigning a constant value to items having an occurrence frequency in the community item popularity data less than a selected frequency and a value that reduces with an occurrence frequency higher than the selected frequency.

41. A method as recited in claim 21, wherein
25 a. the item recommendation data relate to one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services, financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-

related goods, travel-related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products, and

5 b. the serendipity-weighted and filtered recommendation output set relates respectively to the one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services, financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-related goods, travel-
10 related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products.

42. A computer-readable program storage device, having a set of program instructions physically embodied thereon, executable by a computer, to perform a method for providing a serendipity-weighted and filtered
15 recommendation, the method comprising:

 a. receiving, by a processing system, applicable data including,

 i. user item preference data and

 ii community item popularity data;

 b. producing, using the processing system, an item
20 recommendation set from the user item preference data;

 c. generating, using the processing system, a set of item serendipity control values in response to the community item popularity data and a serendipity function; and

 c. combining, using the processing system, the item
25 recommendation set and the set of item serendipity control values to produce a serendipity-weighted and filtered item recommendation output set.

43. A device as recited in claim 42, the method further comprising

 a. receiving, by the processing system, the serendipity control function, and

b. generating, by the processing system, the set of item serendipity control values in response to the serendipity control function received by the processing system and the community item popularity data.

44. A device as recited in claim 42, the method further comprising

a. receiving, by the processing system, an input from one of a serendipity filter system administrator and the user relating to the serendipity control function,

b. selecting, by the processing system, the serendipity function in response to the input received from the one of the serendipity filter system administrator and the user, and

c. producing, by the processing system, the set of item serendipity control values in response to the selected serendipity function and the community item popularity data.

45. A device as recited in claim 42, the method further comprising

a. receiving, through an input/output interface, a serendipity function selection control input from one of a recommender system administrator and the user, and

b. selecting, using the processing system, the selectable serendipity function in response to the received serendipity function selection control input.

46. A device as recited in claim 42, the method further comprising receiving, by an input/output interface the applicable data and transmitting the applicable data, by the input/output interface, to the processing system.

47. A device as recited in claim 42, the method further comprising controlling the serendipity function, using the processing system, with feedback data received through an input/output interface, the feedback data including one of a rate at which the user accepts serendipity-weighted and filtered recommended items, a rate at which the user requests additional serendipity-weighted and filtered recommendations, and a rate at which the user accepts non-recommended items.

48. A device as recited in claim 42, the method further comprising
- a. receiving, by the processing system, the serendipity function from the memory system, and
 - b. generating, by the processing system, the set of serendipity control values in response to the community item popularity data and the serendipity function received from the memory system.
49. A device as recited in claim 48, the method further comprising
- a. receiving a request, with the processing system, from the user for serendipity-weighted and filtered recommendations, and
 - b. combining, with the processing system, the item recommendation data with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set in response to the request received from the user.
50. A device as recited in claim 42, the method further comprising
- receiving, with the processing system, user item preference data that includes at least one of unary values, binary values, and numerical values, and
 - producing, with the processing system, the item recommendation set from the at least one of the unary values, binary values and numerical values.
51. A device as recited in claim 42, the method further comprising outputting, using the processing system, the serendipity-weighted and filtered item recommendation output set to the input/output interface.
52. A device as recited in claim 51, the method further comprising transmitting, using an input/output interface, the serendipity-weighted and filtered recommendation output set to a display device and displaying the serendipity-weighted and filtered recommendation output set on the display device.

53. A device as recited in claim 51, the method further comprising transmitting, using the processing system, the serendipity-weighted and filtered recommendation output set to the memory system for storage therein.

5 54. A device as recited in claim 42, the method further comprising receiving, using the processing system, the community item popularity data from the memory system.

55. A device as recited in claim 42, the method further comprising
a) receiving the applicable data, with a single processor of the one or more processors,
10 b) producing the item recommendation set, with the single processor, from the user item preference data,
c) generating the set of item serendipity control values, with the single processor, in response to the serendipity control function and the community item popularity data, and
15 d) combining, with the single processor, the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.

20 56. A device as recited in claim 42, the method further comprising
a) receiving the community item popularity data with a first processor of the one or more processors
b) producing, with the first processor, the set of item serendipity control values in response to the serendipity control function and the community popularity data,
25 c) receiving the user item preference data with a second processor of the one or more processors,
d) producing the item recommendation set, with the second processor, from the user item preference data;

5 e) receiving, with the second processor, the set of item serendipity control values from the first processor, and
f) combining, with the second processor, the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.

10 57. A device as recited in claim 42, the method further comprising
a) receiving the community item popularity data with a first processor of the one or more processors
b) producing, with the first processor, the set of item serendipity control values in response to the serendipity control function and the community popularity data,
c) receiving the user item preference data with a second processor of the one or more processors,
15 d) producing the item recommendation set, with the second processor, from the user item preference data;
e) receiving, with a third processor of the one or more processors, the set of item serendipity control values from the first processor and the item recommendation set from the second
20 processor, and
f) combining, with the third processor, the item recommendation set with the set of item serendipity control values to produce the serendipity-weighted and filtered recommendation output set.

25 58. A device as recited in claim 42, the method further comprising producing, with the processing system, the serendipity-weighted and filtered item recommendation output set under real-time, interactive constraints.

59. A device as recited in claim 42, the method further comprising producing, with the processing system, the serendipity-weighted and filtered

recommendation output set as at least one of unary, unordered recommendations and priority-ordered recommendations.

60. A device as recited in claim 42, the method further comprising receiving, with the processing system, the item recommendation data as at least one
5 of unary values, binary values, and numerical values.

61. A device as recited in claim 42, the method further comprising generating, using the processing system, the set of item serendipity control values in response to the community item popularity data and the serendipity function, where the serendipity function is at least one of:

10 a fixed function having controllable parameters;
a bi-level function that excludes items occurring in the community item popularity data with a frequency greater than a pre-selected upper frequency value;
a bi-level function that excludes items occurring in the
15 community item popularity data with a frequency less than a pre-selected lower frequency value;
a continuous function having a value that reduces with increasing frequency of occurrence in the community item popularity data; and
20 a function assigning a constant value to items having an occurrence frequency in the community item popularity data less than a selected frequency and a value that reduces with an occurrence frequency higher than the selected frequency.

62. A device as recited in claim 42, wherein
25 a. the item recommendation data relate to one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services, financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-

related goods, travel-related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products, and

- 5 b. the serendipity-weighted and filtered recommendation output set relates respectively to the one of musical items, audio/visual items, written publications, articles from written publications, Internet documents, consumable goods, dining and entertainment services, financial service products, real estate, architectural goods, architectural services, automobile-related goods, automobile related services, travel-related goods, travel-
10 related services, images, pictures, works of art, computer-related hardware, computer software and computer-related service products.

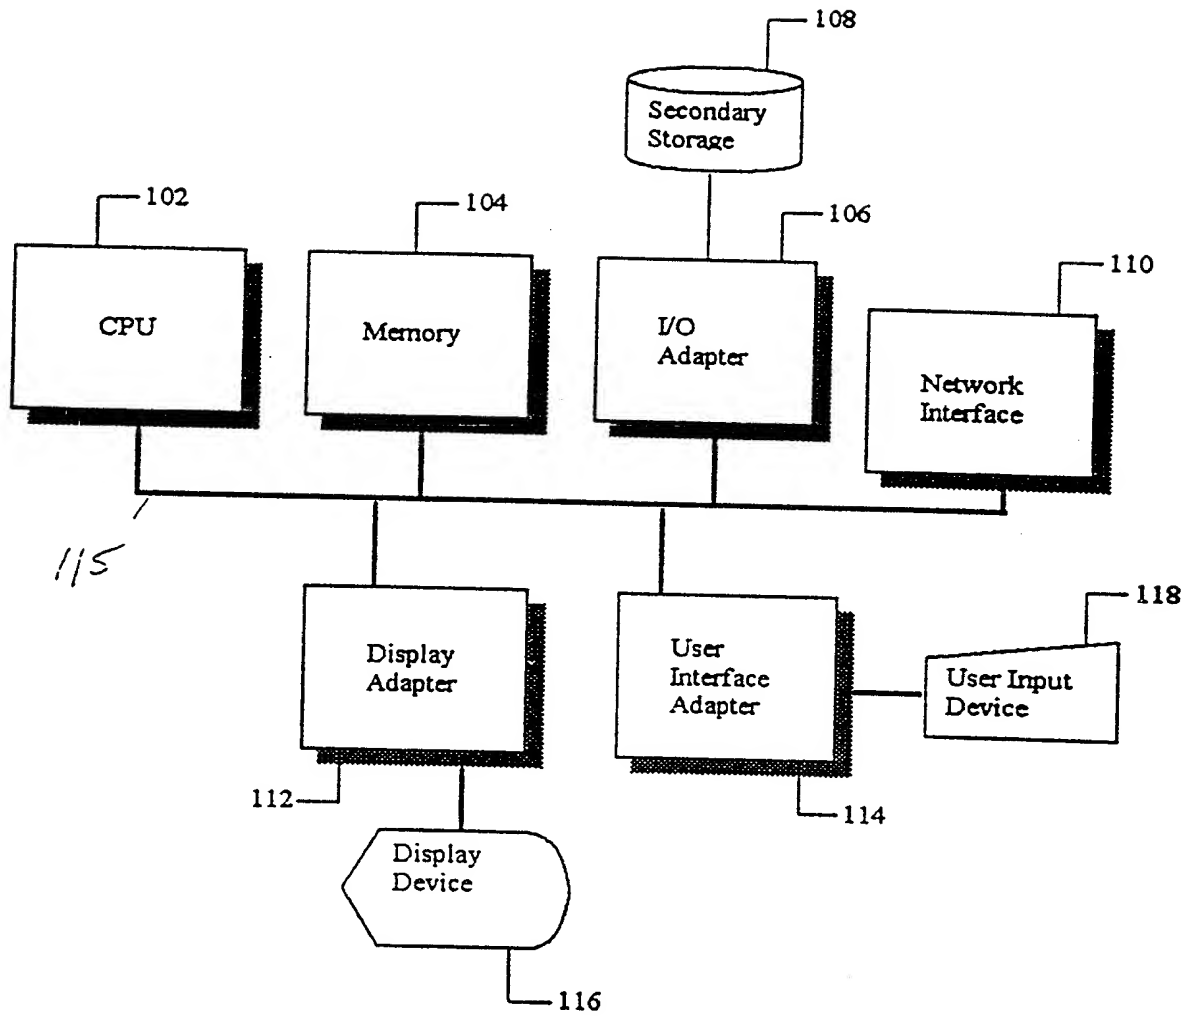
Figure 1

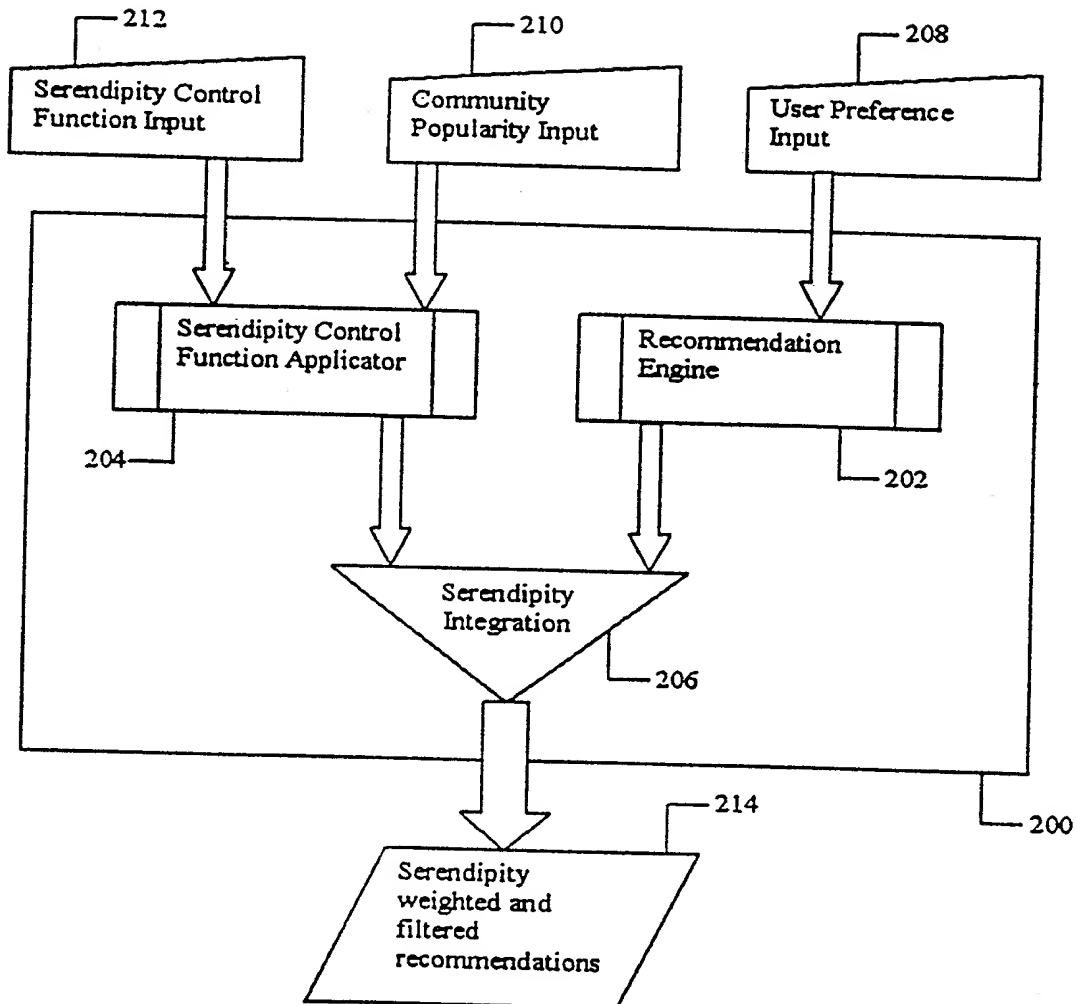
Figure 2

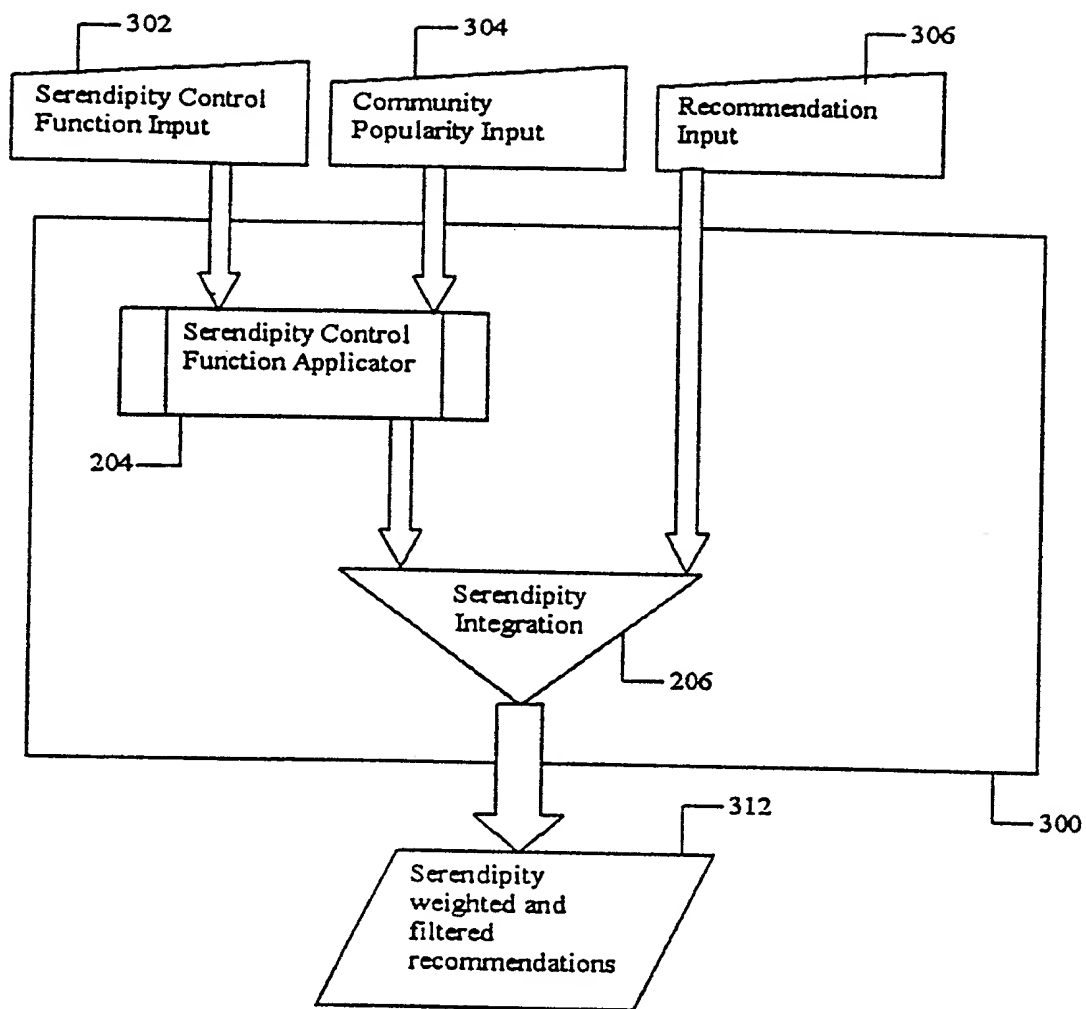

Figure 3

Figure 4

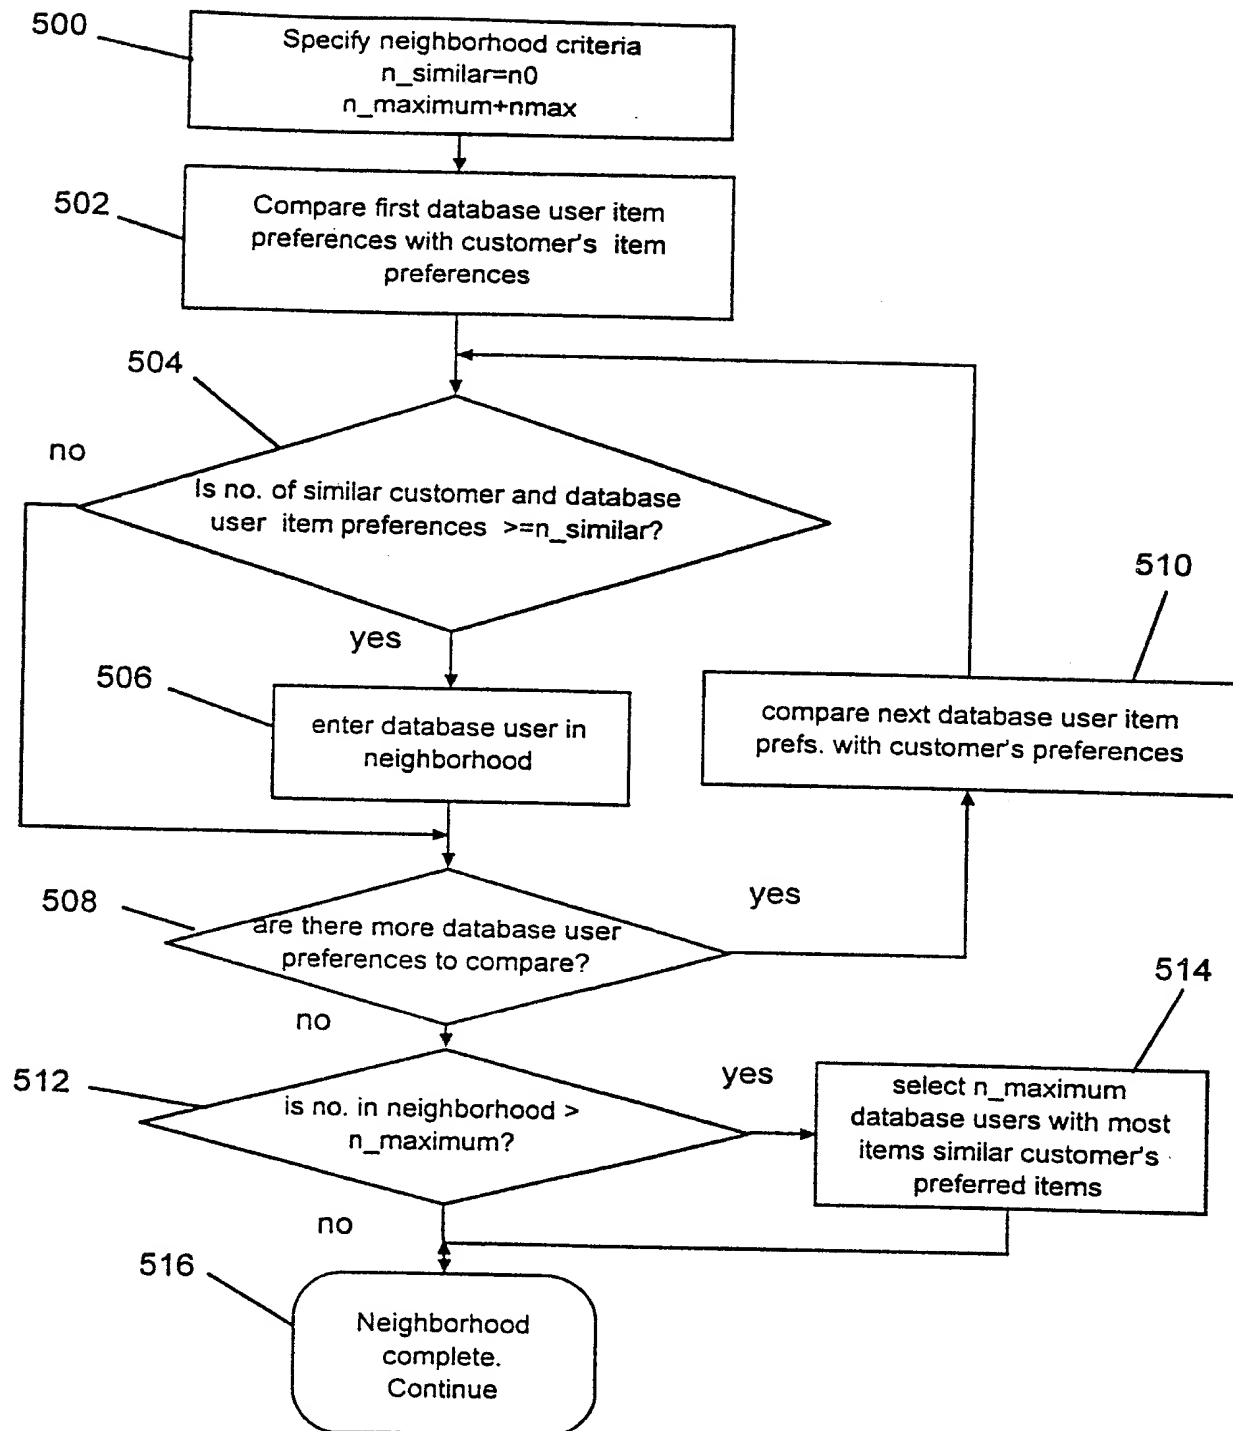

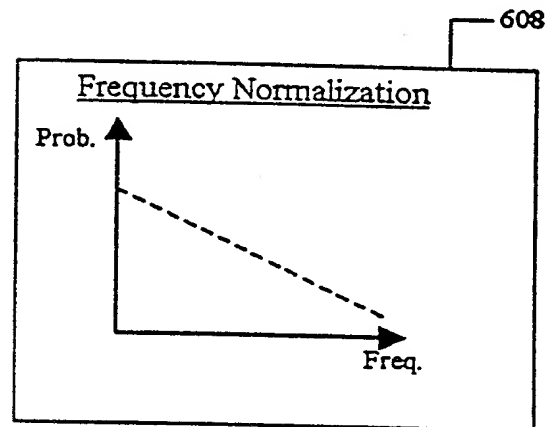





FIG. 5

Figure 6

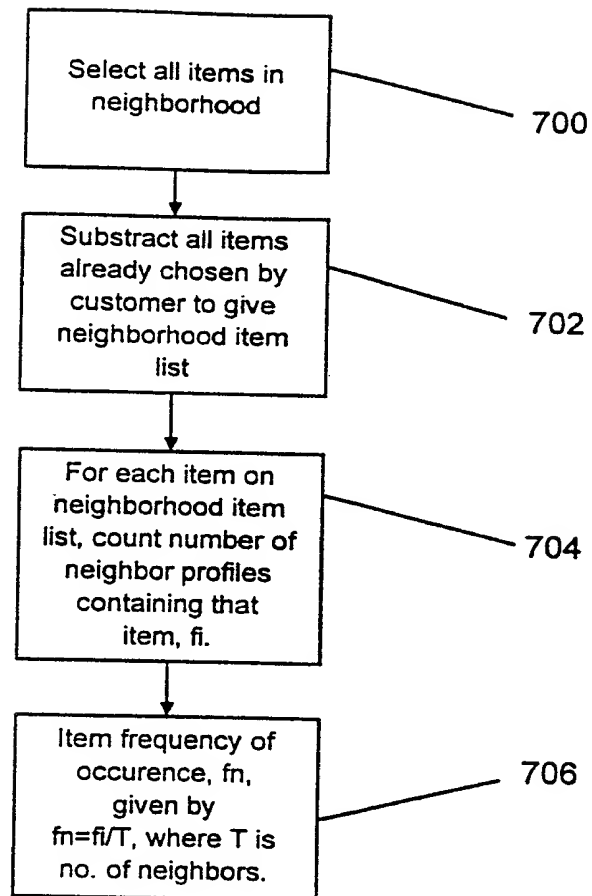


FIG. 7


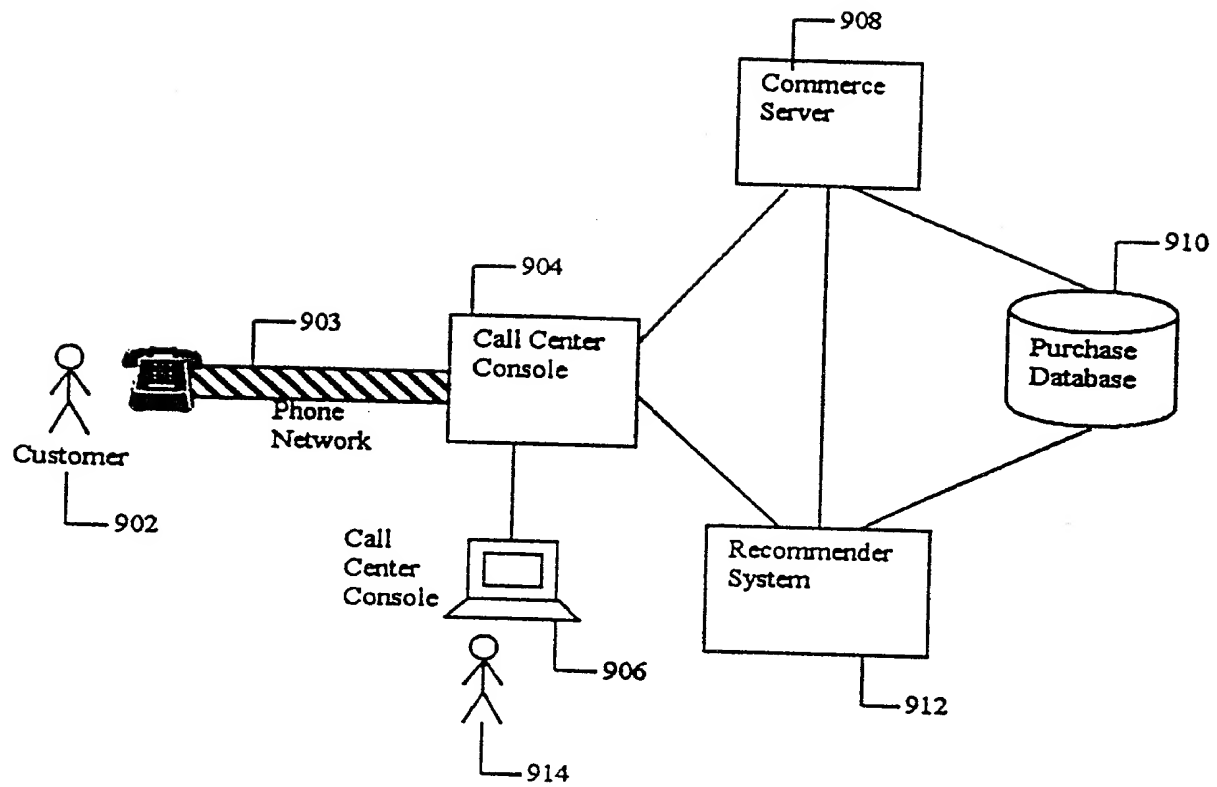
Figure 8

Figure 9

10/17

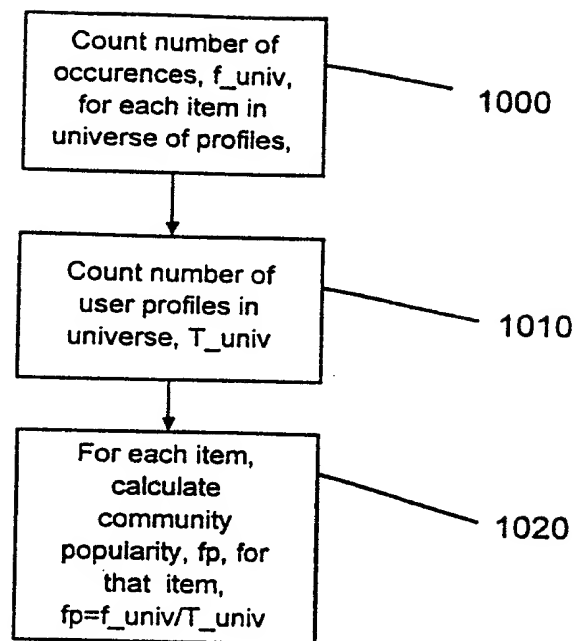


FIG. 10

11/17

Serendipity Control Function Applicator Flowchart 1

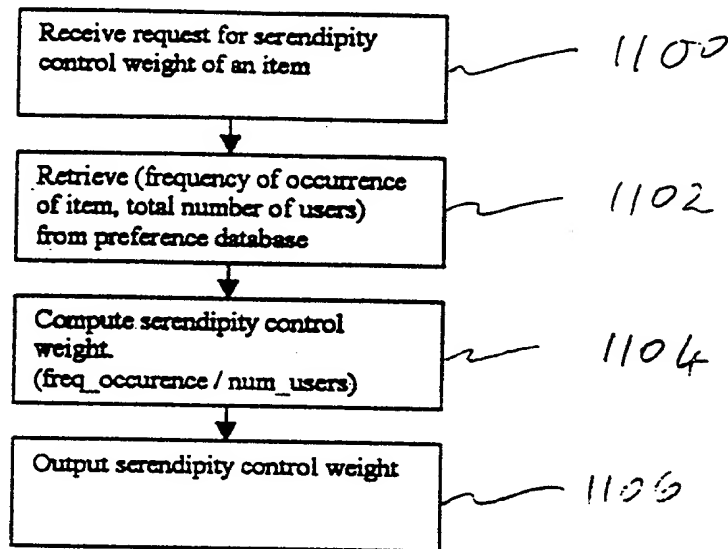


FIG. 11

Serendipity Control Function Applicator Flowchart 2

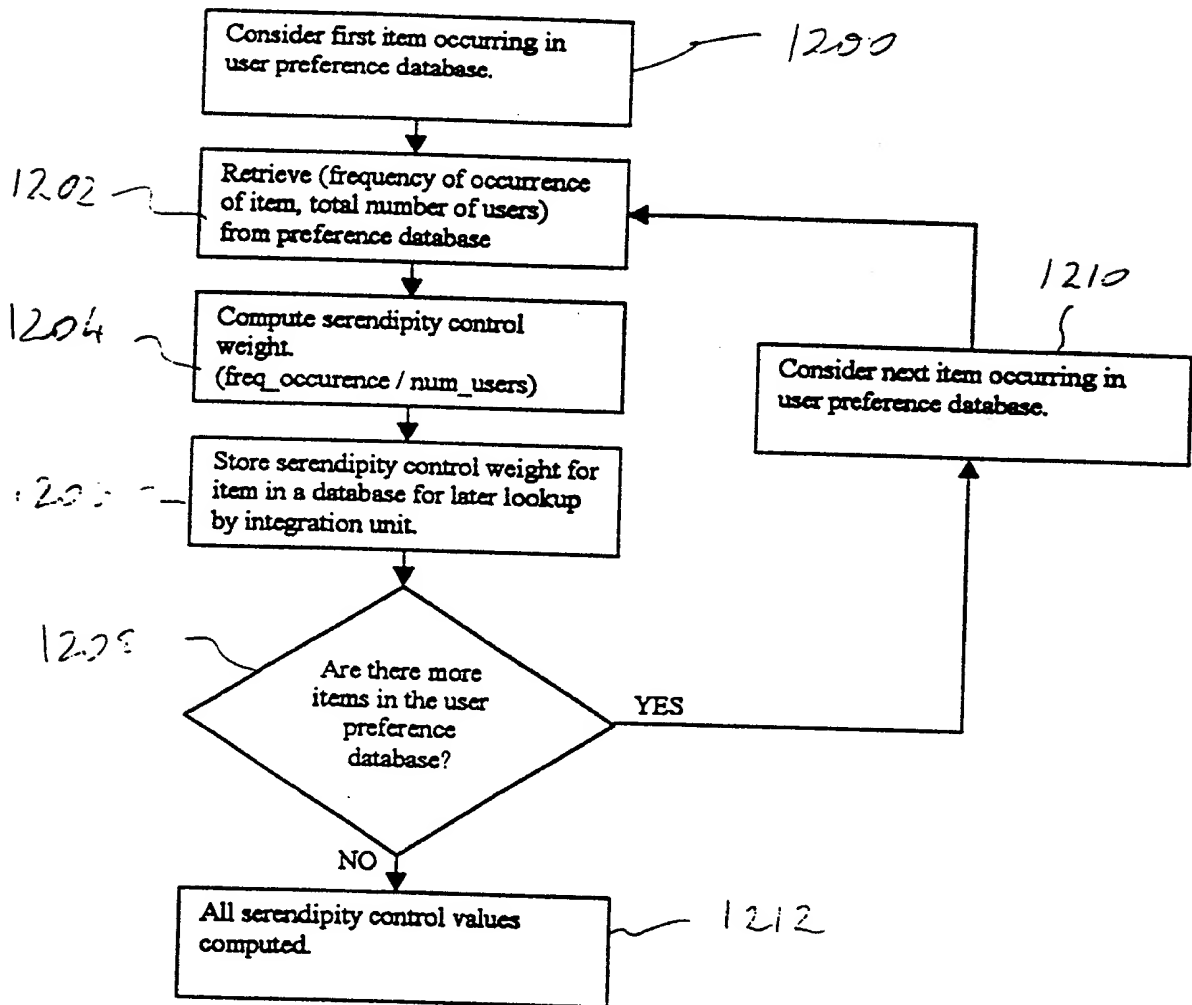


FIG 12.

FIG. 13

Serendipity Integration Unit Flowchart 2

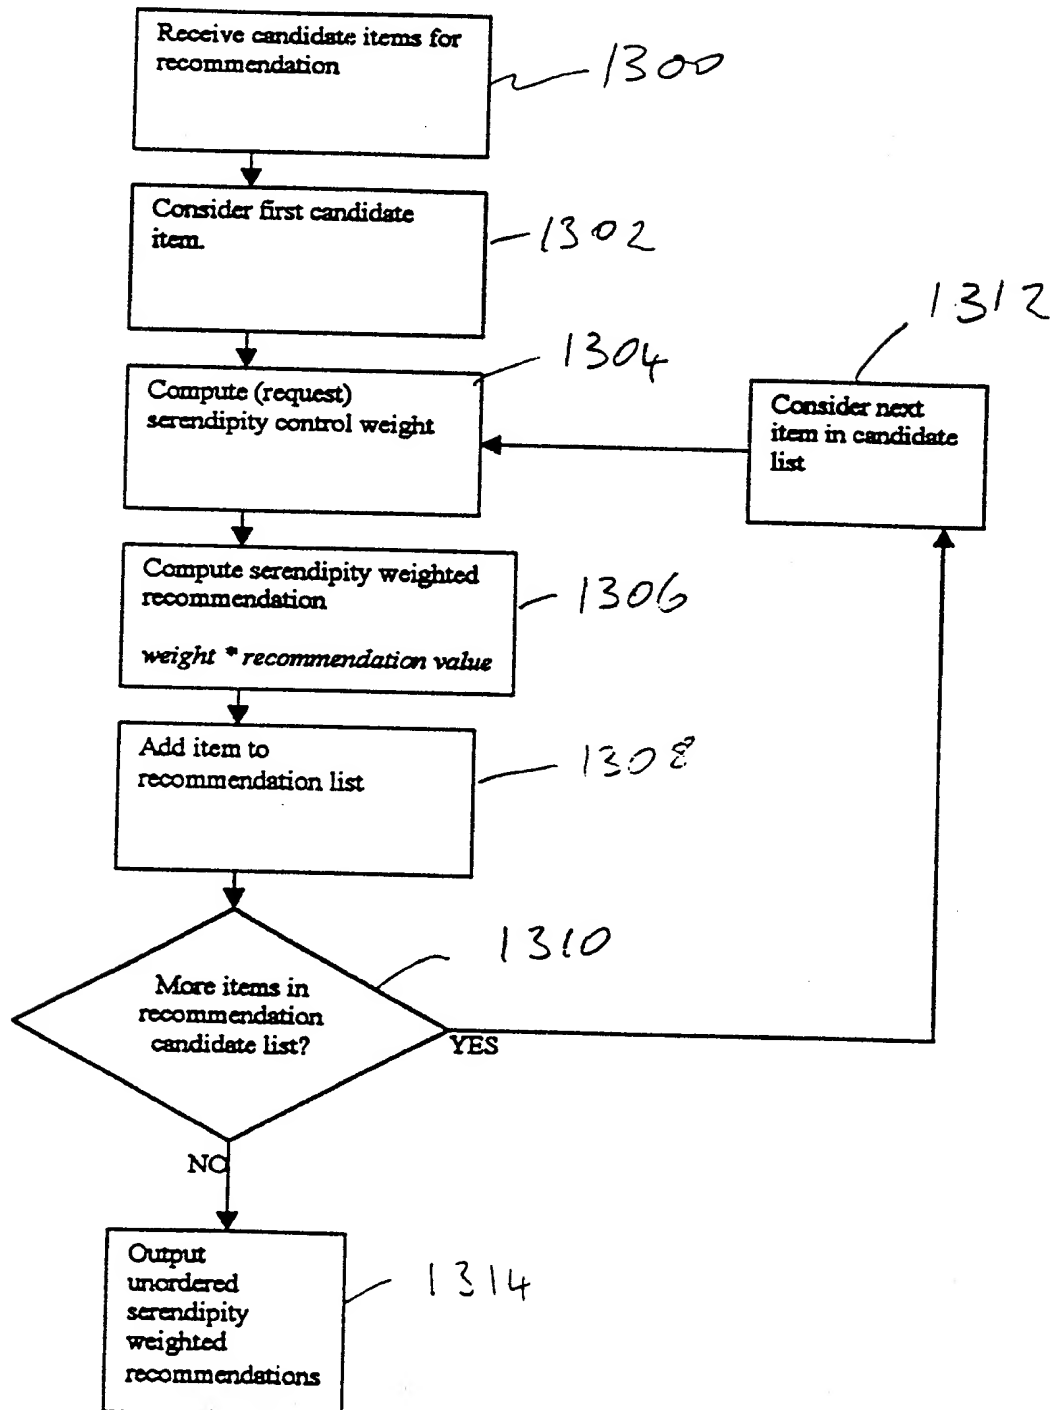


FIG. 14

Serendipity Integration Unit Flowchart 1

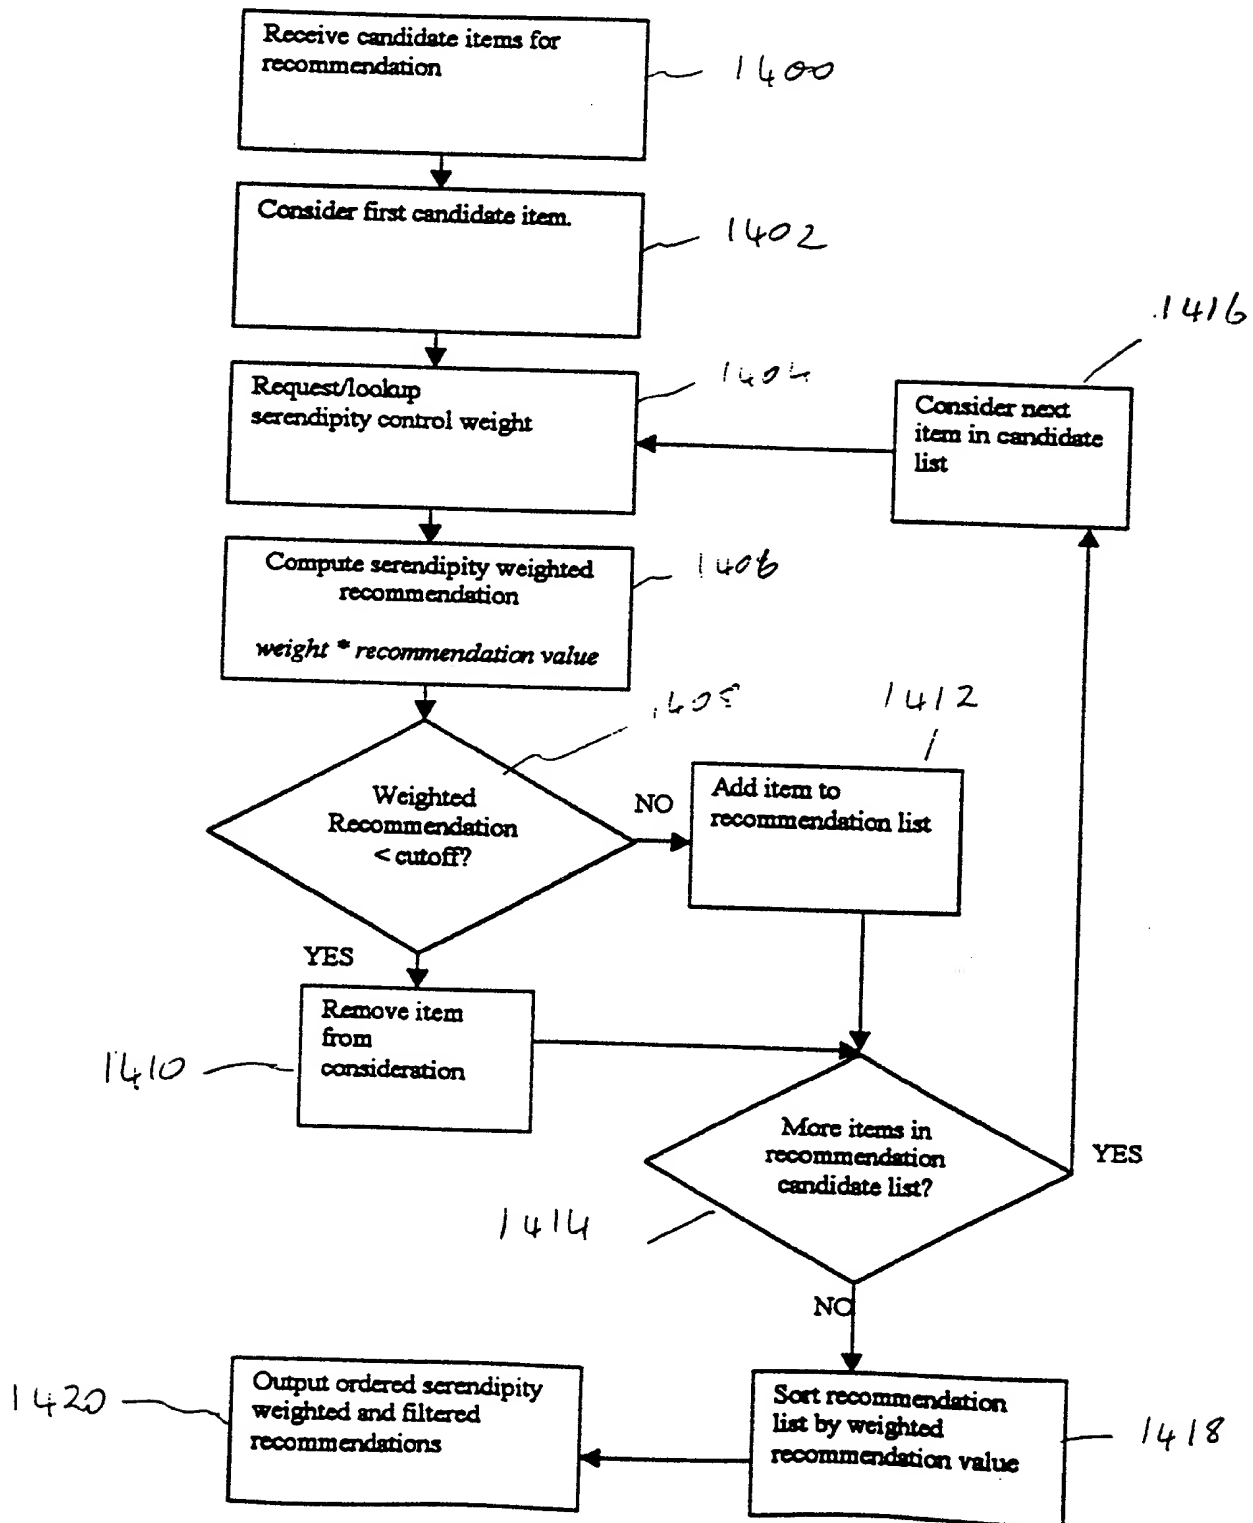


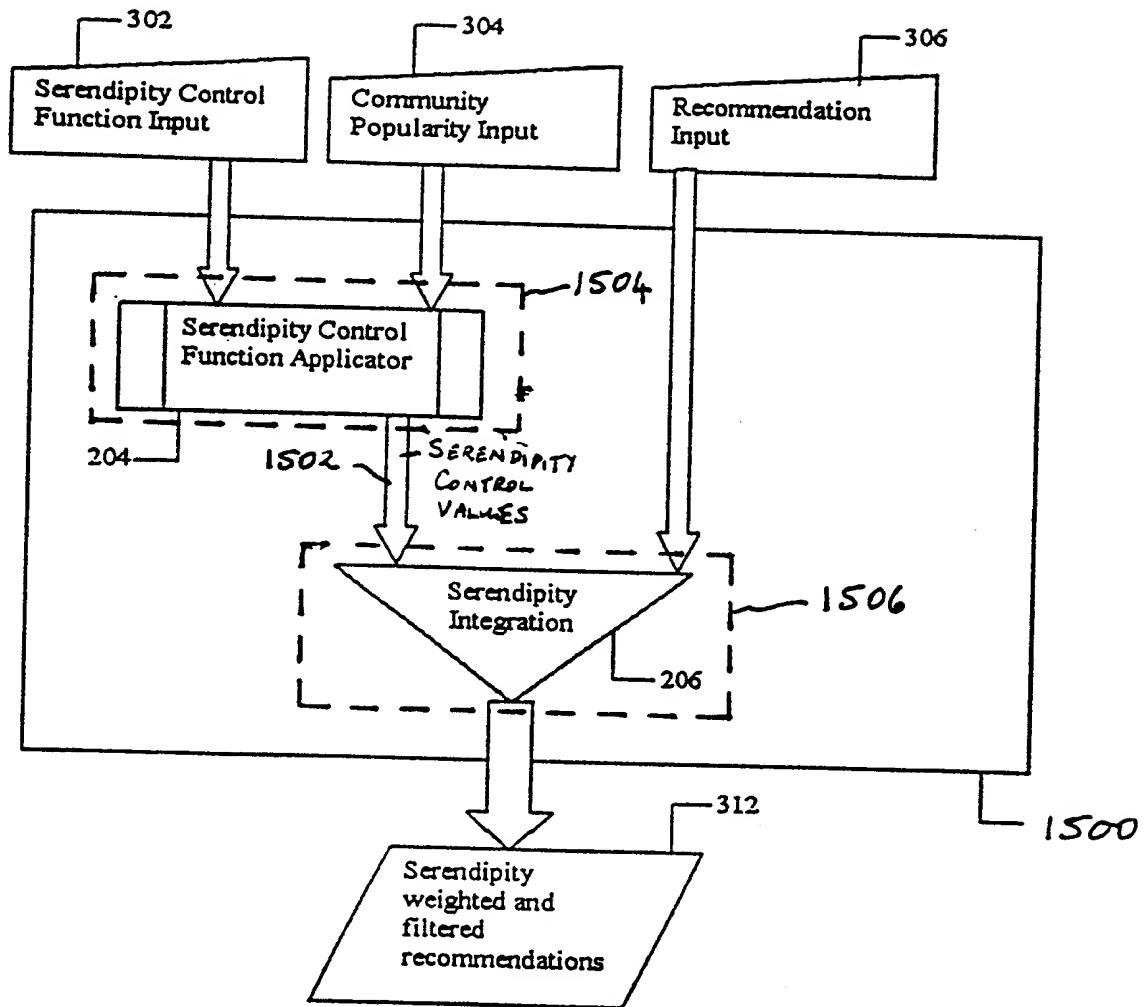
Figure 15

Figure 16A

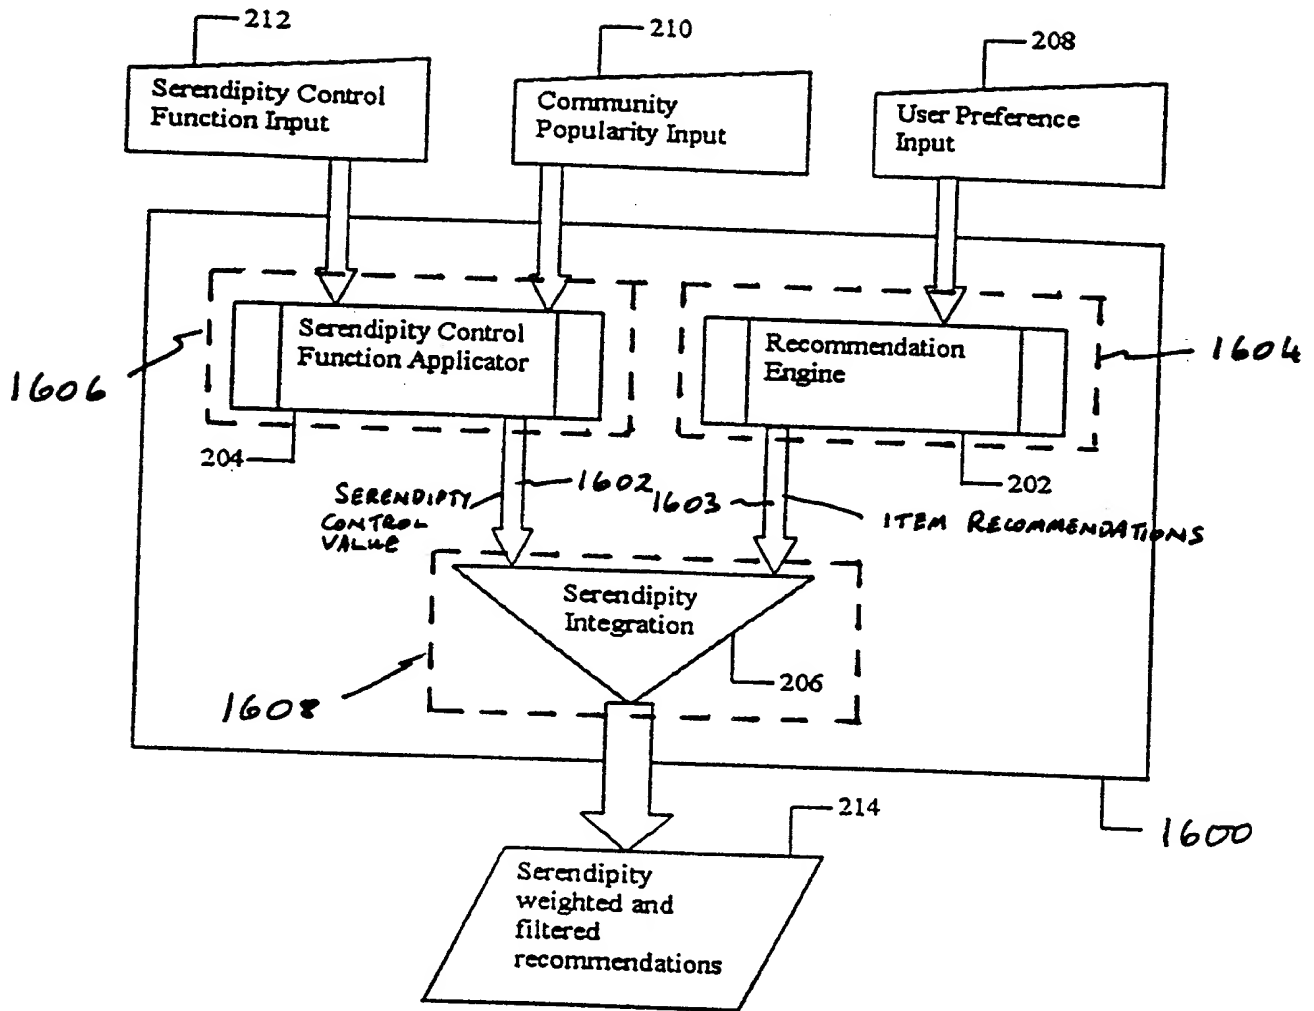
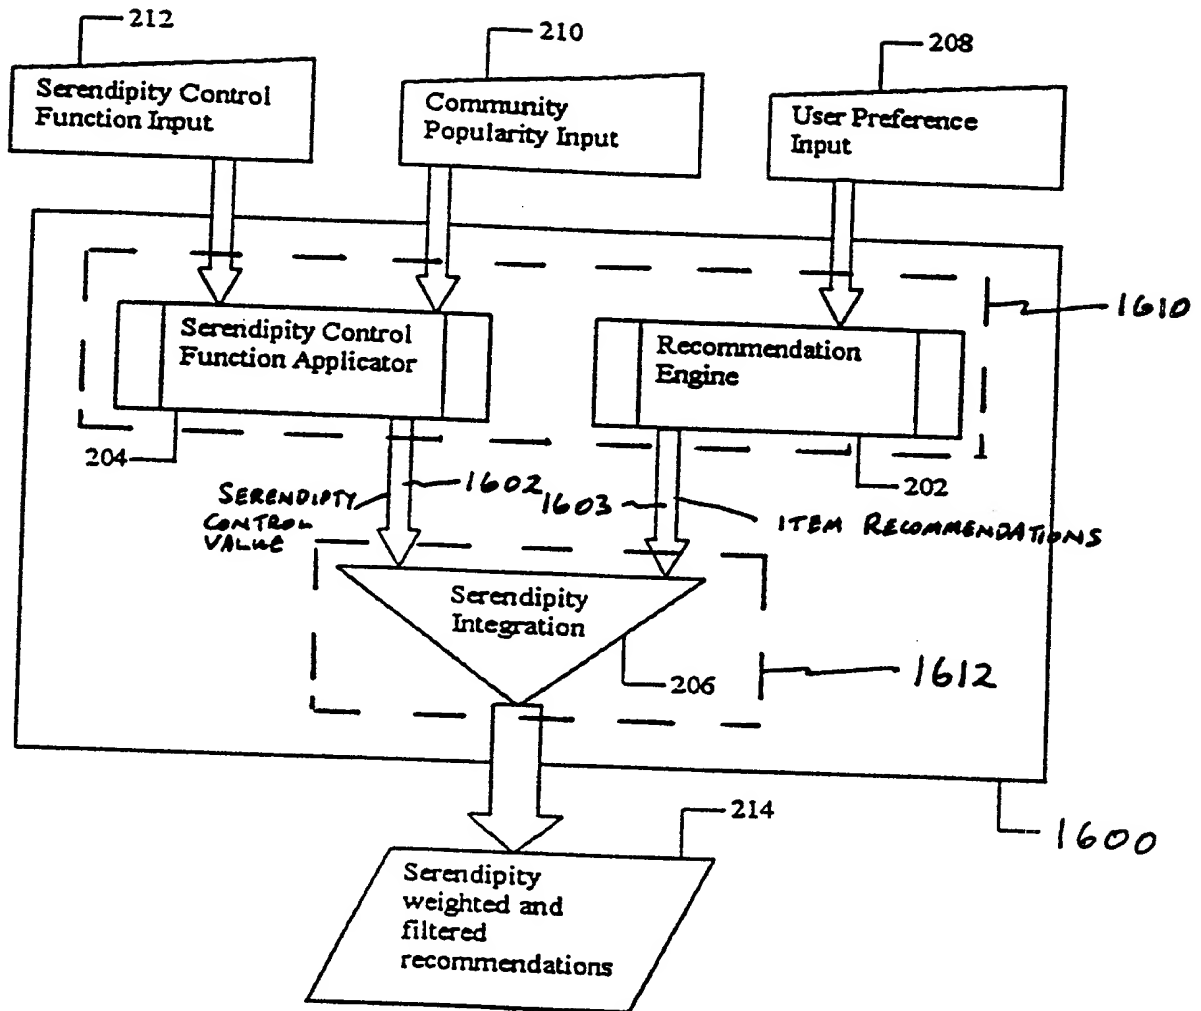


Figure 16B

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/15350

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F17/60 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	LOCKE C: "Intelligent agents create dumb users (?)" ONLINE & CD ROM REVIEW, DEC. 1997, LEARNED INFORMATION EUROPE, UK, vol. 21, no. 6, pages 369-372, XP002121545 ISSN: 0309-314X page 370, column 2, line 6 -page 372, column 1, line 22 ---	1,21,42
A	US 5 704 017 A (HECKERMAN DAVID E ET AL) 30 December 1997 (1997-12-30) abstract column 5, line 25 -column 7, line 55 column 18, line 6 -column 20, line 64 figures 1,2 claims 1-5,19-43 --- -/--	1,21,42



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

4 November 1999

Date of mailing of the international search report

18/11/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Abbing, R

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/15350

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 98 02835 A (MONCHY KATLEAN DE ;POST DAVID A (US)) 22 January 1998 (1998-01-22) abstract page 5, line 5 -page 29, line 28 figures 1-5 ---	1,21,42
A	WO 97 26729 A (ROBINSON GARY B) 24 July 1997 (1997-07-24) abstract ---	1,21,42
A	EP 0 751 471 A (MASSACHUSETTS INST TECHNOLOGY) 2 January 1997 (1997-01-02) abstract -----	1,21,42

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/15350

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5704017 A	30-12-1997	NONE	
WO 9802835 A	22-01-1998	NONE	
WO 9726729 A	24-07-1997	AU 1566597 A US 5918014 A	11-08-1997 29-06-1999
EP 0751471 A	02-01-1997	AU 6282596 A CA 2225790 A JP 11509019 T WO 9702537 A US 5872850 A	05-02-1997 23-01-1997 03-08-1999 23-01-1997 16-02-1999